

Università degli Studi di Palermo
Dottorato in Filosofia del linguaggio e della mente

Between Minds: representing one's own and others'
minds (through explicatures).

Alessandro Capone

Dipartimento di Filosofia, Storia e Critica dei saperi

Settore disciplinare MFIL/05

Il dottorando:

Il Coordinatore:

Il Tutor:

anno accademico 2011-2012

Abstract

In this dissertation, I show how promising pragmatic intrusion based on the notion of explicature can be in dealing with various topics belonging to epistemology and knowledge transmission, such as propositional attitude reports, attitudes ‘de se’, Immunity to Error through Misidentification, knowing how, quotation and indirect reports. The dissertation is divided into two parts. In the first part, I discuss theoretical problems pertaining to the notion of explicatures, the most important of which is ‘Are explicatures cancellable?’ . I argue that they are not. I support this position further through considerations on modularity of mind. I also discuss the picture of inferential pragmatics by pointing out connections between Relevance Theory and Default semantics. I apply notions of modularity of mind and non-cancellability of explicatures to definite descriptions. I also address basic problems pertaining to semantic minimalism and I argue that it is compatible with contextualism. In part 2, I deal with attitude reports and explain how to deal with failure of substitutivity of identicals in terms of pragmatic intrusion. I provide further evidence showing that pragmatic intrusion must be postulated in propositional attitude reports. The discussion on pragmatic intrusion continues with attitudes ‘de se’, which are discussed in terms of semantics and philosophy. I explore various cases of inferences in which ‘de se’ attitudes involve an explicature based on an Ego concept. I argue that Immunity to Error through Misidentification largely depends on semantico/pragmatic considerations and, in particular, on pragmatic intrusion. In this thesis I also explore analogies between propositional attitude reports and indirect reports, from the point of view of substitutivity of identicals. In connection with indirect reports, I expatiate on inferential processes that are non-reflective and independent of context and inferential processes which heavily depend on context (thus the connection between indirect reports and language games). Knowing how is discussed in terms of pragmatic intrusion. I show, in fact, that many problems relating to knowing-how cannot be understood without an application of the notion of explicature. Finally, I provide a view of quotation that is essentially pragmatic and radical, on the grounds that it is more parsimonious than a view which starts with semantics.

Table of contents

Introduction.....	p. 1
Part 1: Pragmatics and modularity of mind.....	p. 11
Chapter 1: Are explicatures cancellable?	p. 12
Chapter 2: Default semantics and the architecture of the mind.....	p. 50
Chapter 3: What can modularity of mind tell us about the semantics/pragmatics debate	p. 81
Chapter 4: The attributive/referential distinction, pragmatics, modularity of mind and modularization.....	p. 109
Chapter 5: Further reflections on semantic minimalism: reply to Wedgwood.....	p. 154
Part 2: Propositional attitudes, knowledge transmission and pragmatics.....	p.199
Chapter 6: Belief reports and pragmatic intrusion.....	p.200
Chapter 7: On the social practice of indirect reports.....	p. 242
Chapter 8: Indirect reports as language games.....	p. 270
Chapter 9: Knowing-how and pragmatic intrusion.....	p.291
Chapter 10: The semantics and pragmatics of attitudes ‘de se’.....	p. 326
Chapter 11: Immunity to error through misidentification (IEM), ‘de se’ and pragmatic intrusion): a linguistic treatment.....	p. 383

Chapter 12: The pragmatics of quotation: explicatures and modularity of mind.....p. 412

Conclusion.....p. 450

Referencesp.454

Acknowledgments

I would like to give my warmest thanks to my tutor Franco Lo Piparo, Marco Carapezza, Francesca Piazza, Marco Mazzone, Luigi Pavone, and Pietro Perconti. A number of other philosophers and linguists should be properly thanked for their comments, encouragement and practical help: Jacob L. Mey, Wayne Davis, Igor Douven, Noel Burton-Roberts, Neil Feit, Istvan Kecskes, Yan Huang, Paul Saka, Igor Douven, Timothy Williamson, Katarzyna Jaszczolt, James Higginbotham, Ernie Lepore, Louise Cummings, Ken Turner, Michel Seymour, Dr. Wedgwood and Keith Allan.

Introduction

In this dissertation I tackled various topics which were interconnected in my mind, proceeding from topic to topic guided by the need to know more and to have access to the mysteries of areas of pragmatics and philosophy of mind which fascinated me. I never got tired to write this dissertation and it was only a temporal limit which made me stop. I am sure that from here I can proceed to more challenging topics, but this will be left for another day. For years, I was fascinated by Gricean pragmatics and in this thesis I was able to use the more conservative as well as the least conservative and innovative strands of research in pragmatics. In particular, I took the notion of (conversational) explicature and from there I moved on to explore topics in the philosophy of mind with particular attention to communicative intentions.

Gricean pragmatics mainly deals with a speaker's communicative intentions as manifested through her speech. Intentions, of course, are in the speaker's head; however, if they were solely in the speaker's head, they could not come out, they could not emerge from her mind and be shared with the co-participants (through communication). This sharing of intentions happens when a speaker manages to get her message across. Now, clearly this practice is not like putting a message into a bottle and putting it to sea in the hope that someone, if anyone at all picks up the bottle, will be induced to open it, read the message and then carry out an action. Sending a message in daily interaction is to shape the message in such a way that the speaker's intention(s) can be recognized by the hearers. Recognition can be aided through contextual cues and clues, consciously used by the speaker and known by the speaker to be conducive to an interpretation process allowing

speaker's intentions and hearer's interpretations (of the message) to match. Cues may lead a hearer to detect an interpretation problem (e.g. a mismatch between the literal and the non-literal level of meaning) and clues can serve to fill the gaps left in the text (Dascal and Weizman 1987). Projection of speaker's intentions occurs through a continuum from a more conscious (and sometimes tactical) process, involving a deliberate planning process determining the inferential procedure and anticipating partially or completely the interpretation moves of the hearer to a less conscious and finally automatic, or even routinized and standardized process. But now it is clear that the pragmatic enterprise is not like bottle throwing in the hope that the message in the bottle will be recovered by someone; the message must be shaped by taking into account the concrete situation of the utterance and the predispositions of the hearers to pick up the cues and clues that direct the communication process. We have a hearer whose history we know, we have a person in front of us, and we use part of the history we share with her and the background assumptions that normally have a bearing on the understanding of a message to anticipate the direction which the interpretation work will take. Anticipating an inferential process is important for the speaker, as he thereby manifests the awareness that his intentions will be most probably interpreted correctly, arriving at a match between the speaker's intentions and the hearer's interpretation. We furthermore use principles of language use (whether we accept Gricean pragmatics or a theory à la Sperber and Wilson) to anticipate the direction which the interpretation process will take. Without a similarity in the cognitive make up of the speaker and the hearer, there would be little chance that what the hearer arrives at are the speaker's intentions. It is the fact that the speaker can rely on this (biologically determined) cognitive similarity and on the awareness of certain principles of cognition, which are valid both for the speaker and the hearer, that guarantees that the inferential path is predictable on an objective basis. So a theory of interpretation, to be of genuine interest, must also be one of the speaker's intentions. Intentions and recovered messages are both the objects of pragmatics, since speakers and hearers both shape their codification and interpretation processes by taking into account principles of language use (whether the Gricean maxims or a single cognitive/communicative principle). It is the intersubjectivity of the interpretation path which Jaszczolt (1999) calls the **social path of interpretation**.

Is it true that semantics only provides schematic information which is then expanded into propositional forms through pragmatic processes? This is clearly one of the problems tackled in this thesis, and one which is likely to receive different answers (depending on the person who answers this theoretical question). My approach is traditional, in accepting that semantics and syntax allow semantic compositionality and form the basis for projecting a minimal level of interpretation corresponding, in many cases, to a minimal truth-conditionally evaluable logical form, which then enters into a process of expansion or completion. I am aware that there are more drastic and radical theories of pragmatic intrusion, but I have my own reservations on the idea that a project that does not take into account the fundamental tension of semantics and pragmatics is really feasible. What is clear is that, whatever our approach to pragmatic intrusion (whether we accept a minimal semantics which is fully propositional or a semantic theory which only provides interpretative schemas) pragmatic intrusion can be put to use in the understanding of philosophical topics if an understanding of the philosophical debate hinges on the meaning and the explicatures of a certain construction – to be called an **intrusive** construction. Which are the topics that can benefit most from pragmatic intrusion? These are clearly those of indirect reports, quotation, mixed quotations, belief reports, knowing how, ‘de se’ attitudes, Immunity to Error through Misidentification, the attributive/referential distinction, etc. So we are sure that by providing multiple connections between linguistic and philosophical topics, we can find out how to fuel the philosophical topics and how to construe them in such a way that we can provide constructive answers.

Is there any hope that we can find out further topics where pragmatics and philosophy interconnect? My answer is positive. At this point we dare say that the issue of how to extend these topics is one of method and not of content. Whenever we find out that a philosophical question can be reduced to a linguistic one (even if partially) and if the linguistic question reduces in one way or another to a pragmatic one, we can be sure that pragmatics can be applied to the resolution (or dissolution) of a given philosophical problem. But how can a philosophical question reduce to a pragma-linguistic question? If a metaphysical or epistemological issue hinges on a logical problem and this logical problem has a linguistic structure which uses connectives or words admitting pragmatic intrusion, then the epistemological issue promises to be reducible to a pragmatic issue.

In general, if pragmatics makes progress and demonstrates that a linguistic phenomenon evinces partial or drastic pragmatic intrusion and if it turns out that an epistemological issue uses that linguistic structure as part of the discussion and the essential part of the discussion hinges on that linguistic structure (or part of the argumentation is based on that linguistic structure), we know in advance that pragmatic intrusion can be applied in a fruitful way to that epistemological problem. The question, essentially, is whether a certain epistemological problem is also partly a semantic problem and whether that semantic problem turns out to be a semantico/pragmatic problem (one allowing or requiring pragmatic intrusion). Of course, it is not easy to assess whether an epistemological problem is also partially a semantic problem and every attempt must be made to disentangle the epistemological from the semantic side of the problem. However, if, despite all attempts, the epistemological problem hinges on some semantic categories (say pronominals or indexicals or some other linguistic structure), and such categories are amenable to pragmatic analysis, then it is clear that applicability of pragmatic categories to epistemology must be granted. Is not this a small triumph for pragmatics? Is not this evidence for the view that much of our philosophy hinges on linguistic structures? I assume positive answers to these questions. This is not surprising since much of our theorizing is done through arguments which are crucially advanced through the linguistic resources of our theoretical arsenal.

These considerations are put to use in my dissertation, where I reflect, first, on pragmatic intrusion and the theory of semantics and then on the application of pragmatic intrusion of the truth-conditional type to epistemological topics, such as belief reports, indirect reports, knowing how, Immunity to error through misidentification and also quotation (quotation being part of a theory of how knowledge is transmitted).

The semantics/pragmatics debate has intrigued me for a number of years. In this dissertation, I have reflected on how to make sense of some initial intuitions about the cancellability of the pragmatic components of explicatures and how to make sense of the issue of propositional attitudes, which surely connects with the issue of explicatures. I have also reflected on what kind of connections could be made and I have ended up connecting the issue of explicatures with the issue of modularity of mind and the issue of propositional attitudes with the issue of indirect reports, ‘de

se' attitudes, knowing-how and quotation. If it is true that the issue of propositional attitudes can be dealt with in terms of pragmatic intrusion, then surely these considerations can be extended 'ipso facto' to indirect reports, which work similarly to propositional attitude ascriptions, to 'de se' attitudes and to knowledge claims (in particular knowing-how). The next step, obviously, was to extend the discussion to knowing how, Immunity to error through misidentification, and quotation.

What have these topics in common? I am mainly interested in a theory of how we can transmit knowledge of what happens in another person's mind – whether in the forms of belief(s) or knowledge. Belief reports, 'de se' belief reports and knowing-how are clearly cases of the mental panorama, they represent the contents of another person's mind which we may try to transmit to another person in the hope that this flux of information **between minds** will be of use to the last person in the transmission chain. Indirect reports and quotation crucially belong to a theory of how the mind communicates information about another's person's mind. When we are confronted with an utterance such as 'Mary said that John is in Paris' even if we are not having direct access to Mary's mind, we have access to it through pragmatic inference. So while, strictly speaking, one need not include indirect reports under the rubric of belief reports, if we give full weight to a pragmatic theory, through an indirect report we are able to reconstruct what bits of information belong to Mary's mind and we know how to separate these from information that belongs to John's mind. As far as I know, the only obstacle against placing indirect reports under the rubric of belief reports is the question of how to separate Mary's from John's mental contents. However, once we can do so by resorting to pragmatic principles and abundant clues and cues, we can be sure that the issue of indirect reports and the issue of belief reports intersect and interconnect at a deep level, to such an extent that it is necessary to examine the two topics together. Since quotation occurs frequently in indirect reports, it goes without saying that quotation also intersects with the issue of belief reports. Surely we can also have sentences such as 'Mary believes that John 'is really crazy'' where quotations are inserted in belief reports and thus the problems of quotation and the problems of belief reports intersect in an interesting way. Belief reports and

knowing how are issues that also intersect in interesting ways. Opacity, pragmatic intrusion, Modularity of Mind, the pragmatics of belief, etc. seem to be applicable to both issues and thus it was useful to apply everything we gained through studying belief reports to knowing how. The connection between the issue of belief reports, ‘de se’ belief reports and Immunity to error through misidentification must also be made explicit. ‘De se’ reports are special types of utterances reporting propositional attitudes (‘believing’ is perhaps the most representative attitude). ‘De se’ attitudes are different in that their representation must incorporate a first-personal component of meaning – which I argued must be supplied through pragmatic intrusion. This first-personal component, I argue, is also responsible for Immunity to Error through misidentification. In other words, when we think utterances such as ‘I remember I was walking in Oxford’ the question of the thinker’s identity (or identification or misidentification) does not arise.

The thesis is divided into two parts. In part one I deal with my considerations on Grice’s circle and the cancellability of explicatures, making sense of them in terms of the theory of modularity of mind. The theory of modularity of mind interconnects with pragmatics and, in particular, with my considerations on Grice’s circle and cancellability of explicatures. My considerations on modularity and pragmatics square well with my considerations on Grice’s circle – and this is no coincidence. It shows that there must be a grain of truth in my theory. Given my interest in this match of intuitions, I thought it could be important to test my predictions on the basis of a discussion of definite descriptions. Readers may be surprised to see that I deal with definite descriptions in terms of modularity – but I think ultimately my intuitions may be of some importance.

In part 2, I deal with propositional attitudes. I claim that belief reports can be made sense of through recourse to the theory of explicatures and I use Relevance Theory to show how these explicatures can be recovered. I was surprised to find out that such considerations could be extended to a theory of indirect reports as well. Then I extended my considerations to the thorny issue of ‘de se’ attitudes, knowing how and quotation. Parts 1 and parts 2 are connected by the idea that certain types of pragmatic intrusion are necessary. The mind is programmed to behave in this way, by calculating explicatures automatically in the default case and by providing

important elements of compositionality through pragmatic intrusion (at the level of Merger Representations (see Jaszczolt 2005)).

It is interesting that I adopted Relevance theory in the attempt to explain how utterance interpretation proceeds and how we interpret minds (in each chapter, I reiterate some Relevance Theoretic notions for the convenience of readers). In my thesis, I adopt a modular story in which the brain has a module, with specialized processes and taking as input only certain types of data, dedicated to theory of Mind. I also adopt the idea that the Principle of Relevance is situated in this theory of Mind Module and can explain much of the procedures occurring in such a module. There is a chapter, in this dissertation in which I claim that this adoption of RT conceptual dowry allows scholars to escape the crucial problems one would have in having to explain otherwise – through Gricean principles which apply to communication – the ascription of propositional attitudes to other's minds through thought (rather than through assertion). Having thoughts also implies using certain modes of presentation of the reference of a certain NP and Gricean pragmatics can explain how modes of presentation are attributed only through more complicated stories. It is clear to me that propositional attitudes is the locus where a theory of cognition and communication diverge most considerably and where a theory of cognition alone is better equipped to handle pragmatic problems. But then I am aware that this is a controversial point – and in any case I am open to the idea that thoughts about others' minds can also be communicated, in which case the Gricean project has no apparent difficulties.

Pragmatics deals with the reconstruction of communicative intentions on the part of hearers and the projection of communicative intentions on the part of speakers. Projection and reconstruction of intentions obey the same principles of pragmatics. Those who say or think that pragmatics should primarily be a theory of understanding and those who say that pragmatics should be a theory of meaning projection are both wrong, unless they make these claims together. It is important that one who speaks should take into account the exigencies of those who hear the message and it is of equal importance that those who hear the message should be bound to the duty of reconstructing the speakers' intentions. The **alignment** between meaning projection and understanding is guaranteed by the fact that

hearers put themselves into the shoes of speakers and speakers put themselves into the shoes of hearers. It is also guaranteed presumably by the psychological reality of cognitive principles, like the Principle of Relevance, which steers communication by balancing cognitive rewards and processing efforts. Obviously, cognitive efforts are justified if they serve to yield appropriate and sufficiently rich cognitive effects. This is enough to orient the communication process in its double dimension of transmission and reception of information (the communicative intention). There should obviously be a sort of correlation between cognitive efforts from the point of view of the speaker and cognitive efforts from the point of view of the hearer; and there should also be a correlation between cognitive effects from the point of view of the speaker and cognitive effects from the point of view of the hearer. In this way, we can recognize the double dimension of transmission and reception of information, justifying the adoption of the same cognitive principles for both processes.

Why is it that Cognitive Principles figure so prominently in explanations of epistemological topics. One reason for this is that epistemology and communication intersect in various and important ways. Epistemology deals with knowledge (basically) and the various ramifications of this concept. And knowledge is not only stored in someone's brain, but it is also transmitted, standardly through communication. Skepticism about the transmission of knowledge is usually unjustified. All we need for believing that we know something by being at the end of an information chain through which knowledge of it has been transmitted is that we have good grounds for believing that our informants are reliable and are also able to select reliable informants as sources of knowledge further up in the transmission chain. So, basically, the agent at the end of a transmission chain will not only deem that the person next in the hierarchy of the transmission chain is a reliable informant but that every person situated in the same transmission chain is a reliable informant and is capable of selecting a reliable informant who himself is able to select another reliable informant, and so forth and so on, until every position in the transmission chain is filled.

Since epistemology and communication clearly intersect, the understanding of many epistemological topics requires a theory of pragmatic inference based on a

general view of cognition and efficiency. We need a theory that explains why inferential processes happen so fast and are often conclusive. Meanings conveyed pragmatically are often fast and automatic. I have argued in several places (e.g. in Capone 2009a) that intentions, once they are projected, cannot be easily cancelled. I have discussed in several places the issue of cancellability of intentions in both conversational implicatures and conversational explicatures, and my conclusion has been that intentions, once they are fixed through the numerous clues and cues of discourse, cannot be undone: hence particularized implicatures and explicatures are not cancellable. This is a bold claim, of course, and it happens that the only ones who put forward this claim are Burton-Roberts and myself. My own claims are furthermore supported by considerations about modularity of mind. People working on modularity of mind (such as Kasher 1991) and pragmatics have long since argued that pragmatic inferences are not modular, because they are cancellable, hence non-mandatory, being the result of the interaction between various sources of information, including beliefs about the world, encyclopedic knowledge, deductive inferences, etc. However, if my considerations about the non-cancellability of explicatures are on the right track, we need to assess the status of modularity of conversational implicatures and explicatures. My own view is that implicatures and explicatures are part of a modular picture of the mind, and exemplify the operation of cognitive principles working within the theory of mind module. In this thesis I argue in favor of the interaction between the Principle of Relevance and considerations put forward by Jaszczolt in her default semantics. In particular I argue that standardization of pragmatic inferences due to the principle of relevance creates links or connections with a default semantics archive. I also discuss, however briefly, the notion of modularization which links default semantics and the principle of relevance. I am not, however, saying that all pragmatic inferences are of the modular type as I concede that at least reflective inferences are of the non-modular type, following ideas by Cummings (2009).

My considerations on propositional attitudes concern belief attitudes, ‘de se’ attitudes and knowing how. In an upshot, my main idea is that pragmatic intrusion adds meaning components to propositional attitudes, such meaning components not being part of their semantics but pragmatic enrichments. While compositionality

cannot be guaranteed for propositional attitudes at the semantic level, it can be guaranteed at the pragmatic level (the level of merger representations, following Jaszczolt 2005). Hence propositional attitudes, 'de se' attitudes and knowing-how are probably the most exciting topics within the theory of mind (serving to justify the idea that compositionality is best done at the level of Merger Representations). I also provide some non-standard considerations concerning Immunity to Error through Misidentification and quotation. I consider the issue of quotation, indirect reports and propositional attitudes connected, since the main problem that emerges from such discussions is how to increment logical forms with modes of presentation of the reference which are provided by pragmatic processes. My point of view is philosophical and cognitive. In other words, I will make sure that considerations based on the Principle of Relevance will intrude into the semantics of propositional attitudes.

Part 1

Pragmatics and Modularity of Mind

Chapter 1

Are explicatures cancellable?

Abstract

The aim of this chapter is to reflect on the necessity of pragmatic development of propositional forms and to arrive at a better understanding of the level of meaning which Sperber and Wilson and Carston famously call ‘explicatures’ (or ‘explicature’) and to support the claim that (the pragmatically conveyed elements of) explicatures are not cancellable – unlike conversational implicatures. While Capone (2003) addressed the issue of the cancellability of explicatures from a merely empirical point of view, here a number of important theoretical questions are raised and discussed. In particular I propose that the analysis of the notion of intentionality and of the nature of pragmatic intrusion will settle the question of the cancellability of explicatures. An explicature can be considered a two-level entity, in that it consists of a logical form and of a pragmatic increment this logical form gives rise to (in the context of utterance). However, both the initial logical form and the pragmatic increment are the target of pragmatic processes, in that we need a pragmatic process to promote the initial logical form to a serious intended interpretation and another pragmatic process to derive further increments starting from this initial logical form as promoted to a serious utterance interpretation.

An explicature is a combination of linguistically encoded and contextually inferred conceptual features. The smaller the relative contribution of the contextual features, the more explicit the explicature will be, and inversely. (Sperber and Wilson 1986, 182).

Introduction

The boundary between semantics and pragmatics has been the object of much recent linguistic theorising and discussions. It is an outcome of dialectical conflicts that better theories emerge, in which a number of errors are purged, arguments are refined and perspectives are broadened. The aim of this chapter is to reflect on the necessity of pragmatic development of propositional forms and to arrive at a better understanding of the level of meaning which Sperber and Wilson and Carston famously call ‘explicatures’ (or ‘explicature’) and to support the claim that (the pragmatically conveyed elements of) explicatures are not cancellable – unlike conversational implicatures¹. This is not a polemical chapter, in fact it is intended to advance the discussion on the semantics/pragmatics debate acknowledging the importance of relevance theorists’ contribution to the issue. Of course, the notion of explicature is originally Sperber and Wilson’s (see Sperber and Wilson 1986), but Carston in a number of articles and in her book has further refined the notion by ample discussion. In fact, this author has written a monumental volume doing much service to the cause of pragmatics, by making it a more rigorous discipline. Yet, if my idea that (the pragmatically conveyed elements of) explicatures are not cancellable is correct, a number of connected ideas in that book must be changed: and it is possible that this theoretical move will precipitate positive consequences on the theory as a whole, as was claimed in Capone (2003; 2006) in a discussion of ‘Grice’s circle’. Here, however, for the sake of simplicity of discussion, I propose to divorce the issue of ‘explicatures’ from that of ‘Grice’s circle’. While Capone (2003) addressed the issue of the cancellability of explicatures from a merely empirical point of view, here a number of important theoretical questions are raised and discussed. In particular I think that the analysis of the notion of intentionality and of

¹ Someone said that, after all, my claim is not original; yet I modestly have to say that I put forward this claim in my 2003 paper, revised and reprinted in 2006.

the nature of pragmatic intrusion will settle the question of the cancellability of explicatures. An explicature can be considered a two-level entity, in that it consists of a logical form and of a pragmatic increment this logical form gives rise to (in the context of utterance). However, both the initial logical form and the pragmatic increment are the target of pragmatic processes, in that we need a pragmatic process to promote the initial logical form to a serious intended interpretation and another pragmatic process to derive further increments starting from this initial logical form as promoted to a serious utterance interpretation.

1. The cancellability of conversational implicatures.

I propose to settle the issue of the cancellability of conversational implicatures before considering the issue of the cancellability of explicatures, as the two issues are connected, due to the fact that through both implicatures and explicatures a speaker intends to have a certain meaning recognized by a recipient. In both cases we are faced with intended messages.

That things are not always easy can be shown by using one of Paul Grice's celebrated examples (1989, 33). A is writing a testimonial about a pupil who is a candidate for a philosophy job, and his letter reads as follows: "Dear Sir, Mr. X's command of English is excellent and his attendance at tutorials has been regular. Yours etc.". Surely the philosopher cannot be unable, through ignorance, to say more, since the man is his pupil; moreover, he knows that more information than this is needed. He must, therefore, be wishing to impart information that he is reluctant to write down. He therefore implicates that he thinks Mr. X is not a good philosopher.

Granting that the teacher manages to convey some message above what is literally said and that the quantity of what is said (as well as the cost of reading all such information) is an element in the interception of the communicated message, I am not sure how the cancellability test applies here. One may contend that the implicature is not cancellable. I doubt that the teacher may write a second letter saying "I apologize for that cryptic message; I was in such a haste; Mr X would have deserved a longer letter, which I now hasten to provide, as follows. (...) In fact, I recommend Mr X for the philosophy job in question". Implicated messages cannot be retracted, in certain **official circumstances**; intentions cannot be 'disintegrated'

(or unimplicated) by further messages if the circumstances are such that these intentions are unequivocally calculated.

An objection may be raised to such considerations. Someone may say that the interpretation of the philosopher's revision is somehow tricky: though the philosopher could deny not having supported his student, he certainly could not assert his support. Implicature cancellations are connected with the capacity to deny having implicated that Q by asserting P, but not the capacity to assert Q as a possible repair to having asserted P. But in a sense, this objection is based on the assumption that the implicature is somehow different from the one noted by Grice, that is to say that the writer is not supporting Mr X. On this interpretation, it makes sense to deny the implicature. But now, assuming that this is a reasonable interpretation, why is it that the philosopher's revision cannot explicitly support Mr X? Presumably the reason is that, in addition to the implicature that the writer was not supporting Mr X's application, there is another implicature salient, namely that the philosopher thinks X is a poor philosopher. This is not an implicature that can be denied. The problem of this example, which is surely tricky, lies in the fact that language is embedded in a social situation, in which rules of conduct partially determine the meanings which words – or their absence - have.

In the light of the following discussion, it may be useful to distinguish here between generalized and particularised conversational implicatures. Generalized conversational implicatures are default inferences, to use an expression dear to Levinson (2000), that is inferential augmentations that get through in a default context, in the absence of particular clues about what the context is like, and in which context solely plays a negative role, in that it can cancel an inference in case a conflict arises between propositions already accepted in it and the default implicature. In the neo-Gricean framework advocated by Levinson, the common ground is thought of as a bucket, that contains all the facts mutually assumed, in virtue of common knowledge or because they have been asserted. The defeasibility of implicature generation is explained in this way: we add the content of a new assertion in the bucket strictly in the following order only if each incrementation is consistent with the contents of the bucket:

a. Entailments;

- b. Quantity Generalized Conversational Implicatures
 - i. Clausal;
 - ii. Scalar;
- c. Manner Generalised conversational implicature
- d. Generalized conversational I-implicatures (due to the principle of minimization). (Levinson 2000, 90).

Particularised conversational implicatures, instead, are inferential augmentations in which contextual assumptions loom large in determining/fixing a communicative intention through reasoning using those assumptions as premises. Of course, there may be disagreement as to the level or degree of conscious reasoning actually occurring in the calculation of particularised conversational implicatures. Relevance theorists, for example, prefer to see these inferences as occurring at a subconscious level. I do not exclude that both modes of inference are available and that we have to distinguish, case by case, between conscious pragmatic reasonings and subconscious pragmatic interpretative processes guided by some pragmatic principle.

Some authors, such as Burton-Roberts, believe that particularized conversational implicatures, in contrast to generalized conversational implicatures, cannot be cancelled, without contradiction of what is intended. This correlates with the intuition that the more evident or manifest a speaker's intention to implicate, the less cancellable the implicature will be. (Burton-Roberts 2006, 10)².

1.1 Why is the issue of implicature cancellation relevant to the issue of explicature cancellation?

The issue of implicature cancellation (in the case of particularized implicatures) is relevant to the issue of explicature cancellation because the cases – at least the most important and delicate ones – of explicatures analysed by Wilson & Sperber (2002) and by Carston (2002a) are cases of inferences in particular contexts. Consider the following utterance produced by Mary in reply to an invitation to join us for dinner. When she says 'I have already eaten' at the level of the propositional form we assign a constituent involving a temporal concept and we understand her intention of saying

² PCI = particularized conversational implicatures; GCI = generalized conversational implicatures.

that she has eaten some food at a time interval ranging from a couple of hours ago to a few minutes ago. We certainly do not interpret Mary as saying that she ate some food three days ago, since this information would no longer count as an explanation for her refusal to have more food. Why is this the case? Intentionality is an important factor in our decision and, in particular, since intentionality cannot be assessed telepathically (to use an expression from Wedgwood (2007)), we must rely on features of the context to point in the direction of the right intention. We discard the hypothesis that Mary means that she ate some food some days ago, as this would hardly be relevant to our invitation. A move that is relevant to an invitation to eat some food is either acceptance or refusal. Since no acceptance has been issued, Mary's utterance can only count as an explanation (of her decision to have no more food) and the explanation, to do the right kind of work, must rely on the assumption that she has eaten immediately before the invitation as human beings are such that they cannot eat food when their stomach is crammed. Try now to nullify this inferential work. I think you will not be able to do so. The reason for that is that inferential communication is due to pragmatics and pragmatics cannot be overruled by pragmatics³. If this were done, then there would be no rational way of communicating with people.

In both the particularised implicatures cases and in the explicature cases, we are faced with the same problem. Pragmatics cannot be overruled by pragmatics. Only semantics can be overruled by pragmatics. In both cases, we are faced with intentionality that is evinced through the pragmatic resources of the language and as such cannot be retracted, as pragmatics cannot be overruled by pragmatics. It is not just the question of intentionality, as Burton-Roberts puts it, that decides the issue. Intentionality, on its own, is inert. It is in the head of the speaker and is not telepathically transmitted to the hearer. It is only through communication that the intentionality of the speaker is inferentially relevant to attributions of intentionality on the part of the hearer and it is inherent in the (actual) pragmatic process that no room is left for cancellability of particularised implicatures or explicatures.

The reason why particularised implicatures cannot be cancelled is the same that prevents explicatures from being cancelled. When a strong intentionality is projected, it can no longer be retracted.

³ I am not, however, saying that a specific pragmatic heuristic principle cannot be overruled by manifest contextual assumptions. I am only saying that the ultimate pragmatic process cannot be undone.

Implicatures, it is generally agreed, only arise if intended and recognised if intended. But then it should be impossible to cancel an implicature: how is it possible to withdraw/cancel what was intended to be implicated and was recognised as intended? An implicature could only be withdrawn/cancelled if it were NOT intended. But then it wouldn't BE an implicature (since implicatures by definition are intended); in other words, there would BE no implicature to cancel. This is relevant to the issue of explicature cancellation only in this: presumably, explicatures must be intended and must be recognised as intended. So the above applies to explicatures as much as to implicatures. And this consideration of regarding explicatures as uncancellable is IN ADDITION to the other considerations which you and I advance for saying that explicatures specifically cannot be cancelled. (Burton-Roberts, personal communication).

Perhaps this is too strong. Should we abandon the cancellability test altogether (thus avoiding considering cancellability the hallmark of conversational implicature)? It is amazing, to say the least, that Grice who seemed happy about coupling his notion of conversational implicature with his notion of intentionality did not notice (or pretended not to notice?) the kind of impasse the two notions were leading to. The only cases of conversational implicatures that are really cancellable are those in which the intention to communicate an assumption *p* is least obvious, evident, and recognizable. Presumably these cases, by strict standards, as Burton-Roberts seems to imply, should not even be categorised as conversational implicatures. If things are so, then it is reasonable to go for the strong hypothesis that conversational implicatures are not cancellable (in practice⁴), which carries with it the implication that explicatures too are not cancellable. Explicatures too, in fact are inferences which must be intended and must be recognised as intended.

Things are not so easy though. We should find it useful to limit ourselves to the claim that only particularized implicatures are not cancellable, involving the recognition of certain intentions (by the speaker) on the part of the hearer. These are

⁴ I remain open to the view that potential implicatures can be cancelled in the sense that their potential is not fully developed in real conversation.

the cases in which the pragmatic resources of the language point towards an unequivocal intention and thus make it awkward or difficult to cancel the intention.

One such case is the one of the attribution of a serious intention to a speaker in virtue of having proffered 'p'. Deciding whether the speaker seriously intended to say (to assert) 'p' is no minor matter and clearly shows the nexus between pragmatic processing and arriving at a non-revisable (non-cancellable) intention. It should be noted, following a point made elsewhere by Bach (2001), that the semantic resources of a language are not sufficient to project intentionality. This is a point reiterated also by Recanati (2007, 37), who says (following Frege) that if a sentence lacks the force of a serious assertion, making the content of the sentence more complex by means of operators such as 'It is true that' or 'I assert that' will not change the situation, since whether or not an utterance is serious is a pragmatic matter – a matter of force and not a matter of content.

Suppose I say 'Bush will be remembered for giving prosperity to USA' (here in the written text there is no disambiguating intonation and it is good that it should be so, because this shows the force of contextual assumptions). Well, although everyone will grant that I am using the conventions of English in expressing my thought, it is not clear that literal meaning is a firm guide to the speaker's intentions. You will not just look at the words I used, but you will most certainly embed the sentence in the context in which it was produced. Given the mismatch between the semantic content of the sentence and the propositions true in context (it is enough to look at the current disastrous state of the stock exchange), you will use your reasoning abilities to decide that my intention really cannot be that I want to say in a serious way that Bush brought prosperity to USA. Now, Bach's considerations are nice because they imply that no semantic device whatsoever can settle the issue of the seriousness of the speaker's intentions (this is a point he does not press). Suppose I say: 'Look, I am seriously telling you that Bush brought prosperity to USA'. Surely it could be useful to use the words 'seriously' to indicate that I am speaking with serious intentions. The hearer will give the speaker another chance and will try to find out again whether such intentions are plausible and match with the propositions usually accepted in the broad context of the conversation. Yet, after the search has produced no context in which the intentions could be serious, the hearer will give up such a search and will settle for an interpretation in which the speaker could not be really serious. The truth is that if literal meanings are not sufficient to guarantee that the

communicative intention behind the sentence is serious, no use of literal devices for expressing a serious intention will do. They will be undermined by the pragmatic resources which augment the semantic resources of a language. If we were zealous enough we should be able to describe a great number of inferences of this kind, in which the contextual clues are rich enough to allow the hearer to recover an unequivocal intention. Suppose that my sister asks me whether I would like to eat more chocolates. Yet, I detect from the quality of her voice that she is not sincere. Here we are before a case in which the words say one thing and the facial expression and the quality of her voice say quite another. Yet the intention is in the speaker's head. It is neither in her words nor in my head. This is a case in which it is hard to prove what my sister's intentions are. Yet, if she uses some conventional resources and she uses appropriate contextualization clues we can have access to her professed intentions, at least. The analogy with particularized implicatures here is that we have a message and we want to recover the intention behind the message (the intention in the head of the speaker) on the basis of contextual clues. As in the case of particularized implicatures, the speaker could try to deny having had (or not had) a certain intention, but she is betrayed by her tone of voice or by contextualization clues and, thus, it is possible to fix her intention even if the speaker thinks that she could evade the question of what her intention really is.

In a sense, it is a good thing that the task of rendering an intention serious (of allowing hearers to say that a professed intention is serious) is to be assigned to pragmatics (as Bach 2001 says) and not to semantics. Semantics can be overruled by pragmatics (as far as the assignment of serious intentions is concerned), but it is not the case that (viceversa) pragmatics is overruled by semantics. Furthermore, at least in principle it should never be the case that pragmatics overrules pragmatics. It is a logical impossibility that pragmatics should be overruled by pragmatics. While it is plausible that semantics is overruled by pragmatics, as far as the issue of determining the seriousness of the speaker's intention is concerned (and it is plausible because there is the residual possibility of determining whether a speaker's intentions were serious or not), it is a logical impossibility that pragmatics should be overruled by pragmatics because, having eliminated the possibility that semantics could determine whether the speaker's intentions were serious or not, now there is no further possibility that the seriousness of the speaker's intentions be expressed at all. By contemplating the possibility that pragmatics be overruled by pragmatics, we are

now facing the possibility of living in a world in which everything goes and in which chaos and luck prevail. I think everyone should reject the hypothesis that pragmatics is overruled by pragmatics – and this leads to the idea that particularized inferences are not cancellable. If they were cancellable, then that would amount to accepting that pragmatics can be overruled by pragmatics. And this is hardly a palatable view, as we have seen.

Now if these considerations are plausible, it is clear that the matter of inferring the speaker's intentions given a context of utterance and the clues which are used by the speaker (knowing that such a use will lead the hearer to infer a message) is such that it is not easy at all to retract intentions. As an example, consider the case of a public person (a famous and important politician) who harasses his female secretary by making an indecent proposal. Suppose he uses a line of defence according to which he did not really intend his message. After all, if what Bach says is true and we need pragmatic processing in ascertaining whether someone seriously intended to assert 'p', this line of defence would not be totally unreasonable. Here a problem opposite to the one noted by Bach arises. The problem is not so much how to prove that the speaker's intentions were serious, but whether one could, instead, prove in court that one's intentions were not serious. Could not the politician say that he was just joking? The problem has to do with what kind of interpretation an utterance is given in a context, given that it is clear to the speaker (or at least it should be clear to him) that the contextual resources of the language will lead a hearer to infer a certain intention and that the context in which the utterance was produced will make it impossible to deny having had that intention. Pragmaticians have not given enough thought to this kind of problem. They take for granted that conversational implicatures (or explicatures) are cancellable, but in fact, in a number of contexts, it is not possible to lead the hearer to infer a certain intention allowing her the freedom to cancel that inference. The fact is that, as was pointed in Capone (2005a), we should give attention to a number of linguistic phenomena in which the individual intentionality is expressed by resorting to the social intentionality involved in performing a certain action. In Capone (2005a) I claimed that a number of inferences arise in context and are simply not cancellable. They are not cancellable because there are conventions or rules of semantic interpretation that are based on discourse practices and on contextual clues. These pragmatic phenomena are inferential

contributions on a par with conversational implicatures, yet being based on actual aspects of context, they have all the features of particularised implicatures, and, hence, cannot be cancelled (at least we predict that it is not easy to cancel them). Consider utterances such as “I saw you” in the classroom. The teacher notices that Michelangelo (his favourite student) whispers the answer to a question to his desk mate. The teacher says “I saw you”. This is not just an accusation, but an order to Michelangelo to stop what he is doing. How can this speech act be transformed into the speech act “stop prompting”? It is the social situation with the rules and expectations governing students’ obligations and teachers’ tasks that promote the **inhibitive interpretation** of “I saw you”. In this context, it is out of the question that the utterance could count as a compliment – such an interpretation simply cannot occur. In fact, no matter how highly the teacher thinks of Michelangelo (maybe even admiring him for wanting to help his fellow students), and even though Michelangelo knows that the teacher has this positive opinion about him, it is unlikely that he will choose the tortuous path of individual interpretation and proceed from considerations about his teacher’s high esteem for him to the interpretation that the speech act counts as a compliment. Michelangelo will almost certainly prefer the social path of interpretation to his own individual path (Jaszczolt 1999). Thus, he is able to work out that the teacher, despite his high opinion of him, actually wants him to stop whispering answers to his desk mate. This is clearly a case in which, when the communicative intention is fixed, it should not be retracted.

Consider a case in which the individual intentionality and the social intentionality clash. Pippo De Lorenzo wrote a letter to the University Chancellor, protesting against a certain number of injustices after imputing them to the Head of the Faculty. In this letter he threatened to appeal to the judiciary system. A lot of trouble was caused by this letter, as the letter caused strong resentment on the part of the Head of Faculty. After a few years, Pippo De Lorenzo, meeting a mutual friend, tells him:

1) If you meet the head of the faculty, please pass him my regards and tell him that I sincerely wish him well.

On reflection, there is a clash here between the individual path of interpretation and the social path of interpretation (the resulting irony). We do not know whether Pippo

De Lorenzo is sincere in expressing his intentions. However, let us suppose that he is. In this case, the social intentionality prevails over the individual intentionality and no matter how sincere his words can be and how laborious Pippo De Lorenzo's efforts to emphasise his sincerity by using words such as 'sincerely', the conventional effects of the words are undone by the pragmatic inferences that arise in context as a result of the social intentionality. In vain he could add sentences such as the following:

(2) Let him be sure that I mean what I say;

(3) And I really mean what I say.

These could not unimplicate the social intentionality latent in the message in the given context. If he really cares to communicate a serious message of reconciliation, the speaker, who must be aware of the social intentionality latent in the effects of his words in context, should either abort his intentions or carry out some reparatory moves such as the following:

Making peace with the head of the Faculty;

Ensuring that his apologies are accepted;

Creating a cordiality context that justifies his sincere expression of the words he wants to utter through (1).

I assume that some will take the example above as supporting a view of pragmatics based on the hearer's reconstruction of communicative intentions. On the contrary, the example above supports the idea that there can be no communicative project unless the individual and the social intentionality work in tandem. When there is a divorce between the individual and the social intentionality, the communicative project has to be abandoned. Pippo De Lorenzo should, in fact, either abort his communicative project or ensure that the individual and the social intentionality work in tandem, by altering the context through the steps I proposed.

2. Explicatures as developed logical forms

Levinson (1983) opts for a negative definition of pragmatics – pragmatics deals with non-truth conditional meaning. This view is tidy and orderly: semantics is the basis

for conversational implicatures (Levinson accepts the slogan pragmatics = meaning – truth-conditions). However, as a final note, Levinson (1983) voices some doubts that this tidy and simplistic picture can be maintained, mainly due to the examples provided by radical pragmaticists.

Although various authors have talked about the role played by pragmatic inference in constructing a propositional form (e.g. Bach (1994), Levinson (2000), Recanati (2004), Stainton (1994)), in this chapter I shall concentrate on Carston's thought, as crystallized in her 2002 book. Carston's idea of pragmatic contribution to the proposition expressed has something distinctive because, unlike Bach, she believes that pragmatics contributes to what is said and, unlike Levinson (2000), she believes that the inferences that develop logical forms into propositional forms are explicatures, not implicatures⁵. Carston's ideas are similar to Stainton's and Recanati's, but they differ as to detail. The examples that show the role played by pragmatics in fleshing out a propositional form are roughly of the following type:

(4) I am feeling better today;

(5) On the top shelf (uttered by a speaker who realizes that the hearer, making his breakfast, needs the marmalade);

(6) He wasn't wearing his glasses and mistook his wife for a hat-stand;

(7) This fruit is green;

(8) It is raining.

To express a full proposition, (4) must imply a comparison between the present and a previous state (how the speaker was feeling, say, yesterday). (5) is clearly an elliptical utterance: it is not grammatically complete and requires the addition of a subject and of a verb. (6) is a case in which conjunction contributes a causality notion to the full proposition expressed. (7) also needs pragmatic intrusion, as we are left in doubt as to whether the outside of the fruit is green or whether the interior is so (by resorting to contextual knowledge, we may settle the issue); (8) also says that it is raining here, in the location where both the speaker and the hearer are situated.

⁵ It is fair to acknowledge that radical pragmaticists such as Cohen (1971) also discussed the phenomenon of pragmatic intrusion. Yet, I think their contributions were only programmatic, while Carston's contribution to this issue is systematic and fully-developed.

Carston (2002a, 17) believes that examples such as the ones above demonstrate that, in addition to a speaker standardly meaning more or other than she says, the ‘what is said’ of the utterance may itself involve more than the meaning of the linguistic expression used. So she thinks that we have to distinguish two notions: the linguistic meaning, the information encoded in the particular lexical-syntactic form employed, and the thought or proposition which it is being used to express, that is, what is said.

Carston embraces the underdeterminacy thesis, that is the view that the meaning encoded in the linguistic expressions used (the relatively stable meanings in a linguistic system) underdetermines the proposition expressed (what is said). The hearer must resort to pragmatic inference in order to work out the proposition expressed by an utterance.

Carston gives up the Isomorphism Principle, a recent formulation of which is the one in Fodor and Lepore (1998):

If a sentence *S* expresses the proposition *P*, then the syntactic constituents of *S* express the constituents of *P*.

The approach Carston defends allows that pragmatic processes can supply constituents to what is said solely on communicative grounds, without any linguistic pointer, in which case the Isomorphism Principle does not hold. I ought to mention here that Wilson and Sperber (2002) hold an approach to semantic underdetermination similar to Carston’s – but in this chapter I mainly discuss Carston’s idea, as she has been more specifically concerned with explicatures. The move of abandoning the Isomorphism Principle is welcome, because it allows us to assign a proposition constituents which do not appear in the corresponding sentence’s logical form.

3. An alternative view (Bach 1994).

An alternative view of pragmatic intrusion is that by Bach (1994). According to Bach, what is said does not correspond to a full proposition determined through

pragmatic inference, but corresponds to a minimal proposition or to a propositional radical and is constrained by the following assumption:

The elements of what is said must correspond to elements in the linguistic expression (the sentence under consideration).

Bach uses a test for distinguishing what is said; the test is provided by the following schema: ‘S said that...’. He claims that only those elements of the original utterance that can be embedded without infelicity in the schema above are part of what is said.

Apart from his conception of what is said, Bach agrees with Relevance Theorists that pragmatics is needed to flesh out the proposition a speaker intends to express. He calls such pragmatic inferences ‘implicatures’. He mainly distinguishes between two pragmatic processes involved in the working out of implicatures: completion and expansion. Completion is required for those sentences which do not express a full proposition. Expansion is required for those cases in which a sentence does express a full proposition but this cannot be considered to be the proposition a speaker really intends to express.

Bach’s picture is not incoherent, albeit an approach that considers ‘what is said’ to be a propositional level of thought commends itself to a greater extent. Perhaps a testing case for which notion of what is said must be adopted is furnished by ironic utterances. Here, in the absence of rich contextual clues (the original context in which the words were proffered), it would be misleading to quote the words contained in the original utterance, as these may lead to a misrecognition of the communicative intention accompanying them. In any case, I accept that the expression ‘what is said’ may be understood in two senses: what is A-said, the words uttered, and what is B-said, the thought communicated.

3.1. A paradoxical example of explicature

Some authors such as Cappelen and Lepore (2005) are unpersuaded by the standard examples by radical pragmaticists and discuss a specific point on pp. 64-65. They believe that if the standard examples of explicature do not have invariant truth-conditions, it could be shown that no sentence has got invariant truth conditions. Consider the sentence:

(9) John went to the gym.

One could argue that this sentence is not truth-evaluable since one could always go on to ask: how did he go to the gym? Did he walk to the vicinity? What did he do in the gym? Etc. I think that what the authors want to prove is that if we think hard enough, then every example of language use will exhibit semantic underdetermination, simply because we set the standards of truth-evaluability too high.

Montminy (2006, 14) commenting on Cappelen and Lepore's work, writes that C & L's treatment of Incompleteness Arguments conflates *lack of full specificity* with *incompleteness*: it conflates cases in which a sentence is not completely informative with cases in which the standing meaning of a sentence does not determine a complete, truth-evaluable proposition.

The point Montminy probably misses is that there is such a wide gap between the interpretation of (10) as (11) and of (10) as (12)

(10) John went to the gym;

(11) John went to the vicinity of the gym:

(12) John went into (inside) the gym,

that one is tempted to say that this too is a case of semantic underdetermination: the full(er) proposition being provided by enriching the propositional radical (to use Bach's words), which ought to be something like the following (if we accept what Cappelen and Lepore say for the sake of argument):

John went (in)to an area vast enough to include the gym and its close vicinity (say the perimeter of its courtyard).

Surely, (more) fully determining this proposition requires some narrowing down, that is to say the addition of some concept. A move open to Montminy, which he does not make, is to take what Cappelen and Lepore say for good and to argue that this is a case in which pragmatics intervenes to enrich a truth-evaluable proposition; and

there is nothing bad with it. Yet, this would not be like saying that no proposition at all is expressed by (12).

An alternative move would be to deny the acceptability of the data provided by Cappelen and Lepore. In any case, Montminy does not appreciate the real point of Cappelen and Lepore's discussion, which is a refinement of the question: how do we know when something is a full proposition? Is not there some latitude in deciding whether an interpretation (whether semantically or pragmatically accessed) is a full proposition? And could we not push this latitude further up in our search for complete propositions? And, finally, what does the expression 'a complete proposition' mean?⁶ This final question is important, since all researchers in the semantics/pragmatics debate propose the priority of pragmatic inference on the grounds that semantic interpretation does not provide a complete proposition (or a full proposition). Presumably, a full proposition is the minimal proposition that is truth-evaluable. However, if we are pedantic enough, we could always say that a proposition is not truth-evaluable and that we need (further) pragmatic inference to arrive at a truth-evaluable proposition. The problem is: where do we stop? It may be worth noting that truth-conditions mean something different for minimalists and for contextualists. Minimalists are not interested in what the world would have to be like for the sentence/utterance to be true (which they call verification procedure), but merely in the formal procedure: "p" is true iff p, even if p is incomplete. Given this distinction, it goes without saying that contextualists are more exigent when it comes to deciding whether a sentence expresses a full proposition. However, as Lepore (personal communication) says, minimalists are not at all surprised to find out that many of the propositions we communicate are absurd, illogical, a priori falsehoods and they do not think that it is the task of semanticists to account for these uses.

The considerations so far are applicable to Carston's work as well. Is it possible that if one thinks hard enough, every linguistic example requires pragmatic development into a proposition? Carston is not scared of this consequence, as she professes to be interested in knowing whether the gap between linguistic meaning and what is said is a contingent or necessary property of verbal communication

⁶ Burton-Roberts finds that talk of full propositions is bizarre. A proposition, by definition, cannot be non-truth-evaluable. He also asks: Why should a full proposition be the minimal proposition? Well, I agree that something is either a proposition or it isn't and if it must be truth-evaluable. Presumably the expression 'a full proposition' is redundant.

(Carston 2002a, 15) and she has a chapter in which she discusses whether pragmatic intrusion is a necessary feature of human communication. This is not the place for a lengthy discussion of Cappelen and Lepore (2005) – I refer the reader to an interesting discussion of the Inter-Contextual Indirect report test by Corazza (2007, 124; 125). Another place to read about the Radical Contextualist's response to C&L is the paper by Wedgwood (2007) in which he addresses some apparently thorny problems for C & L, mainly the fact that, paradoxically, the semantics advocated by Relevance Theory is more minimal than the one advocated by C&L who allow a limited amount of pragmatic intrusion as far as pronominals and indexicals are concerned; not to mention the fact that the Intercontextual disquotational tests appear not to distinguish between homonyms and between use/mention. The discussion of these problems would take a paper of its own, so it is better to leave the matter here (see also Capone 2008b).

3.2 Explicatures

Explicatures (assumptions which are required to make a proposition truth-evaluable) must be differentiated neatly from implicatures. The notion of 'explicature' is originally due to Sperber and Wilson (1986) according to whom:

- (I) An assumption by an utterance U is *explicit* [hence an 'explicature'] if and only if it is a development of a logical form encoded by U.
- (II) An assumption communicated by U which is not explicit is *implicit* [hence an 'implicature'].

Carston (2002a, 117) argues that, along with pragmatic processes triggered by linguistic expressions, there are 'free' pragmatic processes that determine certain elements of the explicature on a purely contextual/inferential basis. She believes that the content of explicatures comes from two distinct sources, the linguistic expressions used and the context, and it is derived in two distinct ways depending on its source, by linguistic decoding or by pragmatic inference. She claims that the logical form, which is the output of the decoding phase is solely a schema for the inferential construction of fully propositional assumptions.

Burton-Roberts speculates that Carston's theory implies that explicatures are a development of the logical form L of the sentence uttered, if and only if P

(asymmetrically) entails L⁷. For example, if I say “He shrugged and left” meaning (via explicature) “He shrugged and then left”, it must be the case that the latter proposition implies the former (the explicature entails the encoded form it is a development of). Burton-Roberts (2005), however, contends that “If the encoded form can be entailed, it must deliver a truth-evaluable proposition” (p. 397) and this could be a problem with the notion of development. In a later section I notice that cases of loosening like ‘Sicily is a triangle’ cause a problem for the notion that an explicature entails the logical form from which it takes input.

4. Are explicatures cancellable?

Both Burton-Roberts (2006) and Capone (2006) converge towards the idea that explicatures cannot be cancelled. In the following sections, I discuss Burton-Roberts (2006) and Capone (2006), in the hope to extend that discussion.

4.1 Burton-Roberts (2005) on the non-cancellability of explicatures.

It is interesting to see what Burton-Roberts (2005, 400-401) has to say on cancellability (of explicatures) in his review of Carston (2002a). The author claims that it is not possible, in Carston’s own terms, that explicatures should be cancellable. In fact, Carston says that the implicated assumptions that constitute the explicature are part of the proposition expressed and, thus, are truth-conditional in nature. On this view, [+ truth-conditional] does imply [- cancellable]. If none of the truth-conditional content of the explicature can be cancelled, the explicature itself should not be cancellable. Cancellable implicature, then, is a logical impossibility according to Burton-Roberts⁸.

⁷ Burton-Roberts (personal communication) says he is just speculating that Carston, in fact, thinks of explicatures as definable in terms of entailment (A is a development of B iff A entails B). This is a reasonable speculation. Her earlier Principle of Functional Independence had it that A cannot be an implicature of B if A entails B. Since a communicated assumption is EITHER an explicature OR an implicature (for RT), it follows that any communicated assumption that entails the encoded logical form must be an explicature. So, with explicature defined in terms of “development”, it is reasonable to speculate that development should be defined in terms of entailment.

⁸ This reminds us of a worry already expressed in Levinson (2000, 166).

Carston may find the idea that (the pragmatically derived elements of) explicatures cannot be cancelled unpalatable because if her notion of explicature is to focus on the central role of pragmatics in human communication, freezing the implicatures in the notion of non-cancellable explicatures will amount to a non-insignificant concession to truth-conditional semantics. Readers may notice that Burton-Roberts' objections (to Carston) come from the perspective of truth-conditional semantics.

Burton-Roberts considers an example Carston discusses on p. 138:

(13) She's ready but Karen isn't ready to leave for the airport.

Carston says that the explicature of 'She's ready' can be cancelled, because the sentence (13) is not contradictory. Burton-Roberts, in my view, correctly argues that the sentence (13) cannot possibly be contradictory; it is statements that are contradictory: "contradiction must be assessed at the (propositional) level of explicature". Burton-Roberts' position is in line with my own considerations. I also think that the sentence (14)

(14) If the king of France died and France became a republic I would be happy, but if France became a republic and the king of France died, I would be unhappy,

which 'prima facie' appears to be contradictory, is not really so. Here, even though we cannot point to any explicit time variables (e.g. in the form of time adverbials), the possibility of an interpretative ambiguity (in the sense of Jaszczolt 1999) remains open. This is due to the fact that the temporal relations between the constituent sentences of each conjoined complex sentence have not been specified. A contradiction can arise only when we decide on a particular temporal configuration. The evidence of the configuration under which no contradiction arises allows us to say that the sentence is not contradictory per se. It might be objected that my view that 'P and Q' is not a full proposition should be true if 'P and Q' could not receive truth-conditions. But the example clearly shows that 'P and Q' and 'Q and P' do have different propositions and that they are therefore propositions. Fine, but I think we

have to distinguish between the sentential level and the speech act level. At the speech act level, a proposition is richer than at the sentential level. When I say that ‘P and Q’ is not a full proposition (hence there is not contradiction at the sentential level in (14)), I am talking about the sentential level. I could agree in principle that, at the speech act level, ‘P and Q’ and ‘Q and P’ constitute different propositions.

So Burton-Roberts is right in saying that (13) is not evidence in favour of the cancellability of explicatures. In particular, he believes that ‘She is ready’ in (13) can be interpreted in three ways:

(15) Pat is ready at time t to leave for the airport

(16a) Karen is ready at time t to leave for the airport x;

(16b) Karen is ready for something (though we do not know what).

If interpretation (15) holds, the second clause of (13) surely does not contradict it. If interpretation (16b) holds, the second clause of (13) does not cancel it either. So (16a) must be the explicature Carston has in mind. But it is precisely (16a) that is contradicted by the second clause of (13).

Burton-Roberts considers another example discussed by Carston on p. 138:

(17) He ran to the edge of the cliff and jumped.

The explicature of (17) is something like (18):

(18) Lionel ran to the edge of the cliff and jumped over the edge of the cliff.

For Carston the explicature (17) can be explicitly cancelled by saying (19)

(19) He ran to the edge of the cliff and jumped but stayed on the top of the cliff

where ‘jumped’ is understood as ‘jumped up and down’.

According to Burton-Roberts (19) is not an example of explicature cancellation, but a clarification that (18) was not the explicature, in that the transitive meaning

rather than the directional meaning of ‘jump’ is regarded relevant. Apparently, Burton-Roberts says that Carston is not allowed to take an ambiguous verb and then use a sentence that denies one of the two meanings of the verb saying that in this way an explicature is cancelled (see Burton-Roberts 1994 for an interesting paper on ambiguity and explicature; there he argues that ambiguity is not a semantic concept, but at most a phenomenon having to do with utterances; whatever the story we accept, it goes without saying that pragmatics disambiguates utterances).

Now we are faced with an old problem. Sadock (1978) in an influential article showed that cancellation is not a reliable test for conversational implicature, because one sense of an ambiguous expression can be explicitly cancelled in a sentence set up for this purpose. However, if I am right, Carston’s example is surely a case in which a pragmatic inference gets through despite the ambiguity of a lexeme (in this case ‘jump’). So, as Burton-Roberts says, what may appear as something that cancels the explicature in this case results in a clarification move. What is being clarified is that the non-directional meaning of ‘jump’ is promoted through pragmatic inference, which then, as Carston says, contributes to an explicature.

4.2 Further considerations on non-cancellable explicatures

Burton-Roberts (2006) distinguishes between **A-saying** and **B-saying**. A-saying is taken to be the literal words expressed in an utterance, which can be reported in abstraction from the original context in which they were produced (presumably to fix an intention). Roughly, A-saying corresponds to the words actually proffered by a speaker in communication (Burton-Roberts 2005 says that to report what a speaker has A-said we must (and need only) quote her utterance); B-saying, instead, involves the assessment (the individuation) of the thought the speaker explicitly intended to communicate, and this may involve putting together both the words used and pragmatic assumptions of the context to arrive at explicatures and to add these to what was literally expressed. B-saying involves fixing the speaker’s communicative intention. Burton-Roberts (personal communication) adds:

In fact, to report a B-saying you don’t have to use any of the actual words that were A-said. Thus, to accurately report what you B-said when you A-said “Fa caldo” (It’s warm in here) I can report you as having said that it was hot. Similarly a person who

A-says “It’s at 12 o’clock” can be reported by “She said the meeting was at midday”.

Burton-Roberts says that Carston’s notion of explicature reconstructs what is B-said (I construe this as: explicatures correspond to a level of what is B-said). The author says that explicatures cannot be cancelled without contradiction of what is A-said or of what is B-said. They cannot be cancelled without contradiction of what is A-said because what is A-said is what is linguistically encoded and does not yield a truth-evaluable proposition (if there is no proposition at this stage, no proposition can be retracted, or cancelled). Burton-Roberts deduces that ‘cancellation without contradiction’ must mean ‘cancellation without contradiction of what is B-said. But he says:

Then cancellation of explicature is clearly impossible as well. To allow that an explicature is cancellable would be to allow that an explicature can be cancelled without contradicting the explicature (that what is B-said can be cancelled without contradicting what is B-said). This looks straightforwardly contradictory. Furthermore, assuming a normal understanding of what is to be ‘committed’ to a proposition, what it is to ‘overtly endorse’ it and to ‘express commitment’ to it, it is clearly impossible for a speaker to cancel what she has explicated without contradicting herself (Burton-Roberts 2006, 4).

Furthermore, Burton-Roberts notices that Carston’s claim that explicatures are cancellable shifts emphasis from the speaker’s intentions to the hearer’s reconstruction of these intentions, a move that is dubious in his opinion, since both for Grice and Sperber and Wilson (1986; 2005) pragmatics is all about intention.

Burton-Roberts’ insistence on the logical impossibility of cancelling explicatures is something that is immediately appealing. Yet, we have to ponder a bit what it means to endorse or commit oneself to an explicature. Carston says that a speaker endorses explicatures, and that she commits herself to them – yet what is it to endorse a proposition, what is it to express commitment to it? Much depends on the way we define ‘commitment’ and ‘endorsing a proposition’. In a sense a speaker commits herself and endorses a proposition through conversational implicature as

well – and if we go along with what Burton-Roberts says then there is no subtle difference between Particularized implicatures and explicatures (actually Burton-Roberts (personal communication) stresses he only said that with PCI a speaker commits herself to having implicated the proposition). If a proposition is actually implicated, it cannot be un-implicated, that is cancelled without contradiction of the executed intention to implicate). Particularized implicatures are quite strong commitments to a proposition.

However, if we follow Bach's view (Bach 2006a) that all messages express a speaker's commitment through pragmatics, since in any case a hearer must distinguish between a serious and an ironic interpretation of an utterance (see also Lepore and Ludwig 2005), we are led to the view that commitment is really a matter of 'explicature' and then Burton-Roberts is right in his claim that when a speaker commits herself to a certain intention, that intention is no longer retractable.

Burton-Roberts (2006, 5) writes concerning explicature (apparent) cancellation that treating the relevant phenomenon as clarification rather than cancellation seems an obvious solution to an otherwise serious problem of principle with explicature. If one claims that explicatures are cancellable, one must abandon Carston's intuitive account of explicature in terms of expressing commitment to and endorsement of a proposition.

However Carston may reply by distinguishing between explicatures* and explicatures. Explicatures* are just potential explicatures. Explicatures (without the asterisk) are just actual explicatures. This move would presumably follow an idea by Gazdar (1979) for whom potential implicatures' are assigned *automatically* – that is, independently of any actual intention-to-implicate – to linguistic expressions on the basis of their semantic representation. 'Potential implicatures' only become actual implicatures when the relevant expressions are uttered. If they are not consistent with the assumptions available (mutually manifest) in the context, they are cancelled. Inconsistency with the context of utterance – and thus cancellation – means that they cannot have been intended. A potential implicature', is then an implicature that arises independently of an intention occurring in the mind of a speaker.

Using the above ideas, Carston may want to distinguish between potential explicatures and actual explicatures. Explicatures (without the asterisk) are those that

a speaker commits herself to and explicitly endorses. Explicatures* are only potentially endorsable, things which a speaker potentially commits herself to⁹. Thus explicatures* are cancellable (so Carston would say) while explicatures are not. So, in a sense both Carston and Burton-Roberts are right. Yet, I do not despair that (only) Burton-Roberts is right if explicatures are a more restricted class than what Carston takes to be (a move that circumvents some problems noted by Cappelen and Lepore 2005). Suppose that we confine ourselves to calling ‘explicatures’ those inferential increments that are meant to supply a full proposition, where none is supplied by bare semantics, or to rescue a proposition from contradiction or logical impossibility (absurdity). These explicatures are in no obvious way ‘potential explicatures’. They are necessitated by the contingencies of communication and by the fact that logical forms are too fragmentary or present wide lacunae. Since in these cases there are no explicatures*, Burton-Roberts is right in saying that explicatures cannot be cancelled.

What arguments would support the contention that there are just explicatures (actual explicatures) and not explicatures* (potential explicatures)? Burton-Roberts might go back to the definition of explicatures:

D1

An assumption (proposition) communicated by an utterance is an ‘explicature’ of the utterance if and only if it is a DEVELOPMENT of (a) linguistically encoded logical form of the utterance, or of (b) a sentential sub-part of a logical form.

He may say that the distinction between explicatures and explicatures calls for a complication of the definition D1 above, which would have to be changed resulting in something like the following:

D2

An assumption (proposition) (possibly) communicated by an utterance is/would be an ‘explicature*/‘explicature’ of the utterance if and only if it is a (possible) DEVELOPMENT of (a) linguistically encoded logical form of the utterance, or of (b) a sentential sub-part of a logical form.

⁹ I should make it clear that this is not a position Carston has ever embraced. I think that Carston may react to Burton-Roberts in this way, but I have no textual evidence that she may be sympathetic to the hypothetical position expressed in this paper.

There is a further problem. Burton-Roberts says that Carston's notion of explicature reconstructs what is B-said. What is B-said is something that really goes on in the conversation and is not to be described as being potential. If we distinguish between explicatures* and explicatures, we must abandon Burton-Roberts' tidy picture according to which explicatures (unqualified) are a component of what is B-said. Explicatures*, in fact, do not seem to be a component of what is B-said.

Carston may argue that for definitional purposes we should solely use explicatures (without the asterisk), while explicatures* are just a theoretical construct that reminds us of the pragmatic derivation of the inference in question. Explicatures* are only a reminder that some pragmatic processes went on at some point in the utterance interpretation.

The distinction between explicatures* and explicatures is suspicious on independent grounds. Explicatures* have all the properties of Gazdar's potential implicatures, and none of the properties of explicatures as pragmatically constructed propositional forms. To say that explicatures are cancellable amounts to saying that (potential) implicatures are cancellable, and that comes as no news. The fact that there is a stage of pragmatic communication in which an inference is potential (before an explicature proper is calculated in context) is a recognition of the fact that the inference is potentially an implicature. But it is an implicature before it is calculated. After calculation, it becomes an explicature in the right circumstances, being correlated with the speaker's communicative intention. And actual explicatures are necessarily calculated inferences¹⁰. Of course, the only way to test a potential implicature is to check for its consistency with the mutual cognitive environment (the assumptions which are mutually manifest to the participants in the conversation). However, one finds out, by checking whether the implicature is cancellable, that the implicature does not need to be triggered. This is a problem, because all cancellability amounts to in this case is the fact that the implicature is prevented from arising. However, one could also consider the case that the hearer temporarily considers the conversational implicature, but then due to the assumptions that are mutually manifest aborts this implicature. One way to test this intuition is to check with speakers what happens when, say, a scalar item like 'some' is used and then the scalar implicature arising from it is cancelled. Should we say that the implicature is

¹⁰ Specific comments by Burton-Roberts persuaded me that 'explicatures*' cannot but be potential implicatures.

cancelled or that it never arose? Relevance theorists consider scalar implicatures in context, and thus get rid of the problem, as one context may favour a scalar reading, while another may not. However, the problem is what happens when out of the blue one says ‘Some of the students arrived’? Is it the case that there is a scalar implicature or not? I would say that there is. Now, there are two possibilities: a) the scalar item is proffered without a context that serves to abort its scalar implicature, but then the context is extended by a further proposition proffered linguistically that defeats the implicature; this is a straightforward case of implicature cancellation; b) the scalar item is uttered in context, where some assumptions that militate against the scalar implicature are mutually manifest. In this case, the conversational implicature is prevented from arising.

4.3 Capone (2006) on non-cancellable explicatures.

Capone (2006) considers some examples of pragmatic intrusion such as (20), (21) and (22):

(20) If the king of France died and France became a republic, I would be happy, but if France became a republic and the king of France died, I would be unhappy;

(21) Take these three plates to those three people over there;

(22) You will die (said to John who has just cut his arm).

Capone writes that cancelling a causality implicature that allows us to make sense of an otherwise contradictory (or at least highly indeterminate) statement results in an unacceptable utterance: hence in this case it is not possible, in his view, to build the propositional form, while allowing for pragmatic intrusion, and then cancel the related explicature without rendering the discourse incoherent. While in ordinary cases of implicature cancellation, the speaker can still be considered to have said something intelligible, something that is coherent in itself and non-contradictory, in cases where pragmatics contributes in a decisive way to the propositional form, such a contribution cannot be withdrawn without causing havoc.

It is quite easy to show that the explicatures in (21) and (22) cannot be cancelled.

4.4 Further problems for the idea that explicatures are not cancellable

It may be suggested that a crucial way of proving that explicatures are not cancellable is to point out that they are part of a speaker's intentions. Presumably one of Carston's reasons in claiming that explicatures are cancellable is that she thinks the hearer entertains the proposition conveyed by the speaker with a high degree of probability but never with certainty (he can go wrong in the process of utterance interpretation). As Saul (2004) and Burton-Roberts (2005) point out, relevance theorists focus on utterance interpretation, rather than on utterance production, and this may very well lead them away from recognizing the central importance of a speaker's communicative intentions, which must be a guide to utterance interpretation in so far as it manifests itself through semantic clues and pragmatic strategies (see also Bach 1998). Since intentions in some cases are fixed, it goes without saying that explicatures which are the correlate of those intentions, should be non-cancellable. Saul points out that the speaker's intentions (once they are manifested in thought) are fixed and that while the process of interpretation may finally provide one or more interpretations which are or are not in line with the original intentions, the communicative process started with those intentions and it is those which crucially matter. We should not be surprised, therefore, that there are loci in conversation where failure to attribute a certain communicative intention deprives the utterance of truth-evaluability and it is these which make the case of non-cancellability compelling.

If Carston replies that some cases of implicatures too are the correlates of intentions (specifically generalized conversational implicatures) but are nevertheless cancellable (an indication that correlation with an intention does not prove that an explicature cannot be cancelled), one may reply that intentionality comes in various degrees and that we have some weaker and stronger forms of intentionality. Explicatures correlate with the stronger level of intentionality. If there is such a stronger level of intentionality, then explicatures cannot be cancelled, because they express intentions of the strongest type.

Is it reasonable to assume that there are different degrees of intentionality and that, specifically, intentionality comes in the weaker and stronger variety? Although I do not see this point discussed in the literature on communicative intentions in great detail, I think we could appropriate of ideas by Castelfranchi & Paglieri (2007) on a

constructive theory of intentions. As I said, the intentions they discuss are not specifically communicative, they are mainly intentions to act. For example, I may form the intention to send a paper to JL for publication or I may form the intention to write a book. It may be safely assumed that communicative intentions are a subtype of the intentions Castelfranchi & Paglieri discuss, as it is possible for someone to form the intention to communicate P to H to inform her of the truth of P. Telling or informing someone of P can be seen as a form of acting. Well, Castelfranchi & Paglieri provide a detailed analysis of the role of beliefs in goal processing – that is in the cognitive transition from a mere desire (what they call a ‘pro-attitude’) to a proper intention. The main point of the paper is that goal processing and intention revision are largely determined by belief revisions and that, in order to activate, promote, drop, or suspend a goal, an intention, or an intentional action, one has to provide or modify the appropriate beliefs. Now, surely such a constructive theory of intentions presupposes that intentions come in various degrees and concerning the action A, one can have stronger or weaker intentions. Stronger intentionality, as the authors say, is activated by the provision of beliefs – so the greater the connection between an intention and a set of beliefs, the greater the likelihood that such beliefs will have an influence on the intention and will modify its degree of intentionality, by allowing it to move from the sphere of potential intentionality to the sphere of actual intentionality. I assume that as we move from one sphere to the other, one can find different degrees of intentionality. Now, on the assumption that communicative intentions are a subclass of the intentions Castelfranchi and Paglieri discuss – and this assumption is not implausible for a theory in which we stress the connection between language and action – we should suppose that beliefs interact with intentions not only at the level of intention formation, but also at the level of intention reconstruction by the hearer. A process analogous to the intention construction occurs, and the beliefs assumed to be in the head of the speaker, which are mutually manifest to the hearer (because as Austin once said, a man speaks!), can play a role in allowing the hearer to infer the degree of intentionality involved in the communicative action.

Now we return to explicatures. The reason why explicatures correlate with a stronger sort of intentionality is that they arise in those circumstances where there cannot be an ‘out’ for the speaker, where the communicative intention proceeds

along the path of the only intentionality available, outside which all sorts of wild grass grows and imputing different intentions becomes so implausible as to impair rational communication. Explicatures are not there to rescue the utterance from all kinds of defective communicative effects, such as lack of informativeness, lack of relevance or lack of quality, but are there to furnish an uttered proposition, the speaker's thought, in the first place, the condition *sine qua non* for evaluating all other communicative deficiencies. The kind of deficiencies which explicatures have to remedy have to do with the lack of a truth-evaluable proposition or with the lack of a plausible truth-evaluable proposition, one which is not irremediably contaminated by 'a priori' contradiction or logical absurdity. It is exactly these cases which shape intentionality within the strict mould of the rational assessment of the thought the utterance must be taken to express.

On the relevance-theoretic view, we should also consider higher-order explicatures, which are the processes assigning illocutionary force to an utterance. Now, I think that if what we have said about the intentionality of explicatures is true, it must be applicable to higher-order explicatures. I have discussed the issue of the non-cancellability of inferences in discourse at length in my paper 'Pragmemes' (Capone 2005), where I provided various cases of non-cancellable inferences. I would, nevertheless, like to point out, that questions of truth are not at stake in higher-order explicatures. In such cases, therefore, we miss the connectedness between non-cancellability and pragmatic intrusion. We might also have a somewhat different story, considering the fact that utterances can be multi-functional at the level of illocutionary force. I do not expect that all my considerations concerning pragmatic intrusion are applicable to higher-order explicatures.

5. On the connection between non-cancellability of explicatures and pragmatic intrusion.

Suppose it is accepted that explicatures are cancellable. Then we have problems in explaining how it comes about that explicatures are part of the truth-conditional content of a sentence. Surely, one of the advantages of claiming that an inference is truth-conditional is that it is part of the entailments of a sentence (and, thus, non-cancellable). If we do not want the content obtained through pragmatic intrusion to be part of the entailments of the sentence, why should we make so much fuss about pragmatic intrusion at all? After all, we could have a very orderly picture like that of

Cappelen & Lepore (2005) in which what is said is mainly associated with linguistic semantics and the various contextual phenomena which we claim to be part of the truth-conditional content of a sentence could be easily assigned at the level of the utterance. Is this what we want? Presumably a radical contextual claim is one that assigns constituents derived through pragmatics inside the sentential content (of course pragmaticians may differ on specific details of their theories; for example, in Capone (2006) I proposed cases of pragmatic intrusion alternative to the standard ones, which in my view were easily assailable (and in fact were attacked by Cappelen & Lepore 2005; the fact that Lepore who was aware of my examples never replied to them is eloquent)). A story like mine, which considers explicatures non-cancellable is more coherent with the view that pragmatics provides constituents of thought that intrude into the sentential level of meaning, since entailments are not entailments unless one cannot deny them. One of the classical tests for entailments is the one by Strawson: if S entails P, one cannot utter S and say not P. So now my question is, do we have enough courage to argue in favour of full pragmatic intrusion? If we have, then we should accept that explicatures cannot be cancelled.

It may be argued that matters become simpler if we consider the explicature as a two-level entity, consisting of the entailments of the sentence uttered and the pragmatic increment that goes into the explicature. However, this segregational approach will not do, because both the entailments and the pragmatic increments are subject to pragmatic processing. The entailments of the sentence uttered are promoted to being part of the speaker's commitment only after pragmatics rules out, e.g. ironic interpretations. The pragmatic increments are combined with the semantic entailments of the sentence uttered only after such entailments are promoted through pragmatics to intended meaning. At this point, there is no more reason for segregating the constituents of the explicature, the semantic entailments deriving from the sentence and the pragmatic increments based on them and on contextual assumptions.

The objection I expect at this point is the following. But then, if this is the way things are, why also accept that conversational implicatures are non-cancellable? Well, first of all, I only confined myself to the claim that particularised conversational implicatures are not cancellable. Second, the question is not whether

it is advantageous to consider conversational implicatures cancellable, but whether it is fruitful to consider explicatures non-cancellable. The difference in the argumentation in favour of non-cancellable particularised implicatures and of non-cancellable explicatures is that there are theoretical reasons due to pragmatic intrusion for insisting that explicatures are non-cancellable, whereas in the case of particularised conversational implicatures we can only resort to arguments based on intentionality. Thus there are stronger theoretical reasons in the case of explicatures for claiming that they are non-cancellable.

6. Refining the notion of explicature further.

Let us accept the following assumptions by Grundy (2000).

- A. Explicatures amount to constitutive aspects of what is explicitly said;
- B. Explicatures are not linguistically encoded but have to be pragmatically expressed;
- C. Speakers are committed to the explicature of an utterance;
- D. Explicatures are part of what is communicated and, thus, are overtly endorsed by a speaker.
- F. Explicatures are motivated by the indeterminacy of language (see also Grundy 2000).

One should at this point, given all the previous discussion, add that explicatures are inferences that partially use linguistic meaning and partially use contextualization clues in order to fix (determine) the speaker's unequivocal intention. An explicature is, therefore, the reconstruction of an intention on the basis of what a speaker says in communication in response to the need to reach a full proposition. The full proposition reconstructed by the hearer is one that cannot be cancelled explicitly and one that can be distinguished from implicatures due to the fact that implicatures are additional increments on top of explicatures. Explicatures are more fundamental increments, as without them the proposition expressed would be either contradictory or false or communicatively very inefficient even at the sentential level. Explicatures are explicit in that they cannot be denied.

We should distinguish real explicatures from potential explicatures. Potential explicatures, like potential implicatures, are not cases of real intentionality (or of real intentionality assignment) but are hypothetical cases of intentionality assignment. When one deals with potential implicatures and potential explicatures one says something like: On the basis of ‘S’ (the syntactic and semantic features of S) and on the basis of the fact that the speaker had a reason to use S (and did not use an alternative to S), the speaker intends that P, unless F, where F is some proposition to be derived from the assumptions manifest in context. The clause UNLESS F focuses on the hypothetical nature of potential implicatures and *explicatures. However, as far as actual explicatures are concerned there can be no ‘UNLESS F’ clause. The calculation process in actual explicatures is quickly over and nothing else can be done with it. We can no longer think of it and process it further in the light of other assumptions. The process leading to explicatures should be finite (both in the sense that we usually do not spend a long time in coming to an explicature, i.e. we do not ponder on the case of explicature for months, and in the sense that when we acquire the propositional elements that are relevant to the case, we process the utterance, come up with an explicature and close the interpretation process for good). An explicature case is not like a reasoning case where the evidence is sifted and then the experts come to a decision and the decision process could be reopened at any moment (say a trial in which new evidence is crucial and may determine a completely different outcome). An explicature case is closed when the communicative exchange moves beyond the next utterance. Thus, the evidence that is relevant to the communicative process is the evidence available at t , where t is some temporal variable that is indexed to the time of the utterance whose explicature we seek to elucidate and is upper-bounded by t' , where t' is indexed to the utterance next to it. By the time u' is uttered, the explicature of u is calculated on the basis of the evidence available **at the moment** and the case is not opened further. Now, this is important. If we sought further evidence and if u were contextualized say at different moments, t' , t'' , t''' , etc. the explicature could very well be different (given that it arises from the interplay of the linguistically expressed assumption and assumptions available in the context). So an utterance could, in theory, be associated with distinct explicatures. In order to avoid this inconvenience, we have to keep the interpretation process finite. And this is in line with the point that an explicature captures a **unique** intention. If the utterance were considered at different moments,

different interpretative possibilities should arise (given the fact that different evidence might be available at different moments). But this is not possible since we have said that the intention to be assigned must be unique.

Now these considerations are compatible with what Sperber & Wilson (1986) say about utterance interpretation. They also argue that the interpretation process is finite and instantaneous – in other words, when the best interpretation is obtained the process stops. My considerations on the intentionality of explicatures make these considerations cogent and provide further justification for them.

In this chapter it will not be accepted that relevance-theory should just be a theory about the hearer's interpretations. We have seen that the philosopher Saul takes this to be RT's most crucial weakness. Wedgwood (2007) also reiterates that RT is a theory from the hearer's perspective. My considerations presuppose that the speaker's perspective is a crucial element of RT as well. In any case, the main reasons in favour of non-cancellability of explicatures also favour taking RT as a theory of speaker's and hearer's interpretations, a theory which has sufficient power to guarantee that the speaker's intentions and the intentions inferred by the hearer match. They must match because the speaker and the hearer rely on the same principles, because their minds are similar in their constitution and because they share the same cognitive inputs (perceptual stimuli) as well as a number of contextual assumptions that are mutually manifest.

The issue of the non-cancellability of explicatures is inherently connected with the issue of the match between speaker and hearer's (assigned) intentions. It's easy to prove this. If there was, say, no speaker-hearer's match in intentions, we should allow for the possibility that the speaker's and the hearer's (inferred) intentions could go along separate paths. It would thus be possible to cancel at least one type of intentions, albeit not both. But the idea that explicatures are not cancellable simply denies that either one type or the other type of intention can be denied (or cancelled). To accept that explicatures are not cancellable in the sense of accepting that the hearer's intentions are not cancellable (without saying anything about the speaker's intentions) leads to nowhere, since one reason why the hearer's intentions are not cancellable is that they aim to reconstruct the speaker's intentions which are not cancellable. Once we give up the idea that the hearers' intentions aim to reconstruct the speaker's intentions, there is no reason to stick to the idea that explicatures are not cancellable. From the hearer's point of view they could very well be cancellable,

as in the circumstances in which a hearer evaluates different readings of an utterance without being under pressure to uncover the speaker's intentions.

Going back to the issue whether the relationship between an explicature and the logical form giving input to it is one of entailment, one should consider that an explicature consists of both a logical form and of some implicated materials which are added to the logical form (let us suppose that logical conjunction is what serves to connect the two in the most simple cases; instead, subtraction is the logical operation involved in cases of loosening). Now, if this is the notion of explicature we adopt, it is clear that it should entail the logical form it is a development of in the cases where explicature is an operation based on conjunction, since conjunction is truth-preserving and a logical form P surely entails itself under conjunction. Of course, an explicature under this view should be subject to strong constraints: one such constraint is that:

the implicated materials be compatible with the logical form the explicature is a development of.

For an explicature to be non-cancellable we need to assume that both the implicated part and the explicit logical form to be developed are non-cancellable elements of meaning. But in a sense, even the logical form that is being developed is subject to pragmatic inference, since it must be mapped to a serious intention and we have seen that the assignment of serious intentions is primarily a matter of pragmatics. After this is done, on the basis of this pragmatic process of intention assignment, some implicated materials are assigned to the explicature and made compatible with the logical form that gave input to them. It is inevitable that this should be a two stage process, with a double assignment of intentionality.

However, there are pragmatic processes that contribute to explicatures, such as loosening ("Sicily is a triangle") in which one cannot proceed in this way. Surely one cannot map the logical form 'Sicily is a triangle' to a serious intention and the implicated materials 'Sicily has vaguely the shape of a triangle' when they are added to the logical form that gives input to them require a loosening of the intentionality mapped tentatively to the logical form 'Sicily has the shape of a triangle'. Compatibility must be maximized and this is done by loosening one level of

intentionality. Now, we must accept that the explicature must have only one level of intentionality and that conjunction is, therefore, not the right logical operation in the case of explicature derivation in cases of loosening. Explicature derivation sometimes involves addition, sometimes involves subtraction. If addition is involved (ordinary conjunction) then two levels of intentionality clearly merge into one single level of intentionality. If subtraction is the logical operation involved in explicature derivation, then we have two levels of intentionality (non-serious; serious), but only one of them prevails. The implicated materials prevail over the literally stated logical form. In this case, the intentionality of the explicature is inherited from the intentionality of the implicated materials. What I am saying is that explicature derivation requires some (pragmatic) compositionality at the level of the levels of intentionality, due to the assumption (we have accepted) that the two components of the explicature must be compatible. (Pragmatic) compositionality means we must have some principles determining which level of intentionality prevails. In the case of conjunction, compositionality derives from the logical operation ‘logical conjunction’ and the compatibility assumption. In the case of subtraction, compositionality derives from the compatibility assumption plus the assumption that at least one level of intentionality must project at the global level and this is strongest level. The intentionality ‘Sicily has roughly the shape of a triangle’ has a stronger level of intentionality than ‘Sicily is a triangle’ given that, due to shared knowledge, it cannot be the case that Sicily is a perfect canonical triangle. Having decided which is the stronger level of intentionality, this will project at the level of the global explicature. The fact that the explicature, broadly speaking, cannot be cancelled logically implies that in case the two components of the explicature are constituted by two distinct incompatible levels of intentionality, only one can project: they cannot both project as this would jeopardise the notion that the explicature is not cancellable.

It is an obvious consequence of the discussion so far that cases of explicatures based on loosening, in so far as they use the logical operation of subtraction, jeopardise the definition that explicatures entail the logical forms they are developments of. In fact, a loose triangle does not entail a triangle, that is obvious enough. Nevertheless, there should be a logical operation requiring that the explicature entails **something like** the logical form it is a development of. To be

more precise, the explicature ‘Sicily has the rough shape of a triangle’ requires that we look at a triangle to see what the real shape of Sicily is like.

In the light of everything I have said, an explicature is a process of the following kind:

Starting from a logical form S , develop S by bearing the Principle of Relevance into account and add the feature Te (truth-evaluability) to $u(S)$ as a consequence of the consideration that $u(S)/Te$ has greater contextual effects and fewer cognitive costs than $u(S)/\neg Te$.

The approach so far is minimally distinct from Sperber and Wilson or Carston’s in that they argue that in specific cases the search for relevance leads to the construction of explicatures. Instead, I argue on general grounds that explicatures that maximise truth-evaluability are preferable on the grounds of the Principle of Relevance.

I would like to tie the notion of explicature to that of assertion in communication. We will not make much progress in pragmatics unless we recognize that explicatures are part of assertions. The considerations so far significantly cohere with what Stainton (1994, 280)) says about assertions, revising considerations by Sperber and Wilson (1986):

An utterance U is an assertion that P if and only if:

- (a) Either P is the propositional form of U (i.e. P results merely by completing the Logical Form of U – i.e. by disambiguating it, enriching it and assigning it reference) or P could result merely by completing the Logical Form of U and conjoining it with another manifest Logical Form of the appropriate semantic type: and
- (b) P is consistent with the presumption of optimal relevance (i.e. U actually communicates P).

In other words, Stainton too believes that explicatures form part of the asserted proposition and, thus, is implicitly committed to the non-cancellability of explicatures.

Conclusion

This chapter has assumed that pragmatic intrusion is a rather general phenomenon in language use and that Carston's notion of 'explicature' is very important. This notion may need refinement, and in this chapter we have shown what kind of facts have to be taken into consideration for this purpose. Cancellability seems to me to be an important fact leading to some theoretical revision. Furthermore, the fact that in some cases it is difficult to distinguish between implicatures and explicatures if merely empirical facts such as cancellation are considered, will inevitably lead us to tighten up the definition of explicatures.

Jaszczolt (personal communication) says that Carston may find the idea of non-cancellable explicatures problematic in that it goes against the idea of non-inference (context-driven inference) and makes explicatures more akin to unmarked, default meanings – not Levinson's highly cancellable defaults, but certainly Asher and Lascarides (2003) or Jaszczolt's (1999) defaults.

This is not necessarily an implication of what I have written so far. I think that in Capone (2006) I have amply discussed a case of explicature that required some kind of contextual inference. Explicatures are uncancellable not because they necessarily correspond to a level of default reference, but because the purpose they fulfil is such that it makes them uncancellable. If they were easily cancellable, then it would be hard to see what role they could play in establishing the full truth-conditional meaning of an utterance. While it makes sense to say that a potential implicature leaves an 'out' for the speaker, it is not very reasonable to say that explicatures give the speaker an 'out'. The purpose of committing oneself to a proposition is to leave no room for disagreement as to what the speaker actually means.

Chapter 2

Default Semantics and the architecture of the mind

Abstract

In this chapter, I explore the relationship between Relevance Theory and Jaszczolt's Default Semantics, framing this debate within the picture of massive modularity tempered by the idea of brain plasticity (Perkins 2007). While Relevance Theory focuses on processing (see the interplay of cognitive efforts and contextual effects), Default Semantics focuses on types of sources from which addressees draw information and types of processes that interact in providing it. In particular, I argue that Relevance Theory interacts with default semantics by standardizing inferences which are ultimately compressed (to use a term by Kent Bach 1998b) into a default semantics. I briefly discuss potential obstacles to the idea of default semantics coming from the experimental pragmatics literature (e.g. Noveck and Sperber 2007, Breheny et al. 2005) and I support further the idea of the division of labor between Default inferences and the inferences derivable through the Principle of Relevance. In the end, I compare Relevance Theory and Default Semantics, in an attempt to come to a more unified picture.

Introduction

The main topic of this chapter is the relationship between Default Semantics and the Principle of Relevance within the modular architecture of the mind. I propose to come to a better grasp of the interaction in question by utilizing knowledge of the issue of modularity of mind. In particular, I analyze the phenomenon Bach (1998b) calls 'standardization' and propose that once inferences become standardized, they are no longer processed through the Principle of Relevance, given that they can be furnished directly by the Default Semantics archive. I consider potential objections

to this idea, based on experimental pragmatics and arrive at the conclusion that merger representations, which guarantee compositionality at the level of the utterance, take into account both Default inferences and modulated effects due to context. Following Horn (2005), I assume that contextual information may seep into the pragmatic interpretation while the default semantics is considered.

I shall assume the existence of a module called ‘Mind-reading module’ within which the inferential work related to understanding utterances is carried out. I shall argue that the Mind-reading module includes the processes described by Jaszczolt’s (2005) Default Semantics as well as the pragmatic inferential processes for which the principle of Relevance is responsible. I shall argue that Default Semantics describes inferences that make use of inferential shortcuts stored in an archive and also utilizes simple heuristic procedures (fast and frugal procedures in the terms of Gigerenzer 1999). Default Semantics includes the study of the inferential processes explained through the Principle of Relevance and its basic heuristics (e.g. What is not said isn’t) are specializations of the Relevance heuristic. Specializations arise in response to **recurrent** environmental problems and consist of heuristics which, though based on the principle of relevance, have formed a distinct sub-component of the theory of mind, and can be characterized as associative. A certain stimulus will trigger a certain output based on these specialized heuristics as a result of semantic associations.

In this chapter I will not advocate a conflation of the mental and the neuropsychological levels or advocate reductionism of the mental to the neuropsychological level (See Chomsky 2000). A theory about the mind may surely proceed along a separate dimension, resting on deductions on the basis of what we know about language, comprehension and other most important introspection data and analyses of those data (see also Feit and Capone, Forthcoming); however, if the findings about neuropsychology can independently support our speculative considerations about the way language and language use work, this cannot but be a welcome result. Neuropsychological data are normally treated as corroborating evidence, but it is possible even to consider the possibility that the neuropsychological facts may be at odds with theoretical considerations. It is theoretical considerations that matter most. Chomsky is methodologically right: for

example we can make sense of certain patterns of electrical activity in the brain by the notion of semantic deviance. Surely we could not understand what those patterns of electrical activity would signify if we did not have an independent notion of syntactic deviance and of its importance for a theory of language and of the mental. In this chapter, I do my best to disentangle the mental from the neuro-psychological, even if sometimes the two levels will meet.

1. On modularity of mind

In this chapter I presuppose a modular picture of the mind. I shall confine myself to only sketching this picture with broad brush strokes. Work on the modularity of the mind started with Fodor (1983) and was extended in Fodor (2000). Fodor mainly distinguishes between input systems (e.g. perceptual processes) and a central system. While modular processes are encapsulated, domain-specific, specialized, shallow, fast and obligatory, central processes for Fodor are not encapsulated but draw inputs from a variety of domains.

This strict picture of modularity is being replaced by the massive modularity picture, mainly advocated by Carruthers (2006a), Sperber (2005), Wilson (2005), Sperber & Wilson (2002) and Carston (1996, 2002), among others, where the mind consists of myriad modules, each correlating with a certain (specialized) function, whose input domain is highly restricted (e.g. the visual system is for processing visual percepts only, it does not deal with inputs from other organs) and with certain transducers. In this picture modules can share resources, especially if they are situated in close neural areas and if they do not perform concurrent functions. The main advantage of the modularity picture is that it explains how the mind can react in such a fast manner to simultaneous inputs coming from different types of transducers and how it facilitates learning in cases in which two or more different types of input have to be analyzed simultaneously. Another advantage of modularity is that each module, in so far as it is to some extent insulated from the remaining architecture, can be damaged, without affecting the remaining modules. So, breakdown in a module creates limited damage, since other modules are available to process information coming from the outside. People can have their language system damaged while keeping much of the remainder of cognition intact (aphasia); people can lack the ability to reason about mental states while still being capable of much else (autism); people can lose their ability to recognize just human faces; and so forth

and so on. An obvious advantage of massive modularity is that processing of information does not create an informational bottleneck, given that various tasks can be carried out in parallel. Another advantage is that the modular system allows evolution to add new modules or to tinker with existing ones, without affecting the remainder of the system. A modular mind is capable of evolving and adding new specialized mechanisms in response to environmental challenges.

Unlike most other authors on modularity, Karmiloff-Smith (2010) puts forward the view that the mind is gradually modularized during development, as a result of the fact that different areas of the brain are more suitable for dedicated mechanisms (more relevant to certain types of processing). So, on this view, starting out with tiny differences across brain regions in terms of the patterns of connectivity, synaptic density, neuronal type, etc., some areas of the brain are somewhat more suited than others to processing of certain types of input. These ideas are interesting, as we may be open to the fact that modules **emerge**, not as a result of genetic endowment, but as a result of interaction with the environment and with repeatable patterns of experience.

1.1 The mind-reading module.

In this chapter, I am particularly interested in the mind-reading module. Theories of mind have largely been of two types: the theory-theory of the mind and the simulation approach. The theory-theory of the mind sees mind-reading as essentially scientific thought that is based on generalizations. To make an example, if you know that human beings suffer severely due to the loss of their dear relatives, then by seeing a person who has lost her father you make the prediction that that person is severely suffering (of course, like scientific theories, this sort of mind-reading can go wrong in circumstances in which people deviate strikingly from certain (near)-universal dispositions). The most influential theories of this type were worked out by Wimmer & Perner (1983), who studied how children's minds developmentally change in relation to false-belief tasks.

The simulation approaches see mind reading as running a simulation. You put yourself into another person's shoes and create certain conditions in imagination (suppose I am terminally ill; what would I do?) and then run a simulation on those conditions. This would be unlike theory-theory approaches, in that all you do is to experience your own reactions in response to simulated conditions. However, you

have to quarantine those states of mind which are likely to interfere with the simulation, by NOT letting in states of mind which belong to you and not to the person whose behavior you are trying to simulate.

In this chapter, in consonance with Sperber (2005), Wilson (2005), Sperber & Wilson (2002) and Carston (1996, 2002), I shall try to explain mind-reading in connection with pragmatic inferences (inferences arising from verbal behavior) by accepting the existence of the Principle of Relevance that processes inputs in a very fast way and provides shallow inferences, that is to say inferences triggered by procedures that are designed to provide responses to environmental stimuli in real time, without getting bogged down in laborious reasoning. These inferential processes are fast, automatic, and schematic – in other words they follow what Gigerenzer et al. (1999) call **fast and frugal heuristics**. In connection with this approach, Carruthers (2006) writes:

Most cognitive scientists now think that the processing rules deployed in the human mind have been designed to be good enough, not to be optimal. Given the speed of processing is always one constraint for organisms that may need to think and act swiftly in order to survive, evolution will have led to compromises on the question of reliability. Indeed, it will favor a *satisficing* strategy, rather than an optimal one. (Carruthers 2006, 54).

I produce below simple examples of the type of fast-and-frugal heuristics proposed by Gigerenzer et al. Suppose you are required to answer the question: which of two cities is the larger. A simple heuristic can be useful – you can choose the only one of two options that you recognize.

Another heuristic to be used in response to the same question can be the following: you may look first at beliefs about which properties of cities have correlated best with size in the past. If having a top-division team correlated with greater size in the past, then you select the town which has a top-division team, while neglecting the town which does not have one. If none of the two towns has a top-division team, you move on to the next best predictor of size.

The nice thing about these fast-and-frugal heuristics is that they come with stopping rules; you know when you can stop and you do not indefinitely process information in order to arrive at an optimal answer. What is good enough will suffice, and you stop there.

Before closing this section, I need to remind readers of the kind of evidence generally bearing on the question whether there is (or not) a mind-reading module and whether certain mind-reading tasks are executed within it. The most widely cited study showing that there is a mind-reading module is Baron-Cohen et al.'s (1985) study of **autistic children**. Autistic children are impaired in mind-reading tasks, and they seem to be blind to the notion that other people have minds. Thus, an autistic child will drag his father as if he were a toy. Autistic children, however, seem to be impaired only in the kind of tasks that involve attributing intentions to others, but they need not be impaired in more generic cognitive tasks.

That mind-reading activities may involve automatic, fast heuristics is proven by patients afflicted with **Williams' syndrome**. This disorder results in an average IQ of around 50, combined with linguistic abilities and social skills. People with Williams' syndrome have good abilities for mind-reading and communication but poor general reasoning abilities. This dissociation seems to support the existence, within the mind-reading module, of a sub-module (or more sub-modules) dedicated to fast and automatic inferences concerning a speaker's intended meaning. (See, however, Perkins 2007, for a deeper discussion).

The existence of this sub-module dedicated to mind-reading is supported by another dissociation. Patients with **Asperger's syndrome** have good general reasoning abilities but serious impediments in mind-reading. These people can use general reasoning to compensate for the lacking special-purpose skills. (See Wilson 2005 for a deeper treatment of this point).

For lack of space I cannot expand on these ideas, but I finally refer the reader to Happè, Loth's (2002) important paper on the dissociation between the mind-reading module in connection with actions and the mind-reading module in connection with communication. This is evidence that the Relevance theory module is a sub-module of the more general mind-reading module.

2. Default Semantics

Now that I have discussed the issue of modularity of mind, I turn to Default Semantics, as my aim in this chapter is to situate Default Semantics in a modular view of the mind, which helps explain the relationship with the principle of relevance.

I discuss the important approach to inferential pragmatics developed by Jaszczolt in a series of publications starting from her 1997 seminal paper (Jaszczolt 1999, 2005, 2009a), which she called ‘Default Semantics’. In discussing it, I propose crucial modifications of the framework.

A crucial feature of this framework is the centrality of the notion of the **speaker’s intention**. Linguistic actions – like non-linguistic ones – are animated by intentions and are successful in so far as hearers recognize them.

The theory elaborated by Jaszczolt is complex and multi-faceted. However, here I freely draw on the central ideas of her theory. Jaszczolt (1999) notes that there is a tension between the individual and the social path of interpretation and she correctly remarks that it is the social path of interpretation that must win. The individual path of interpretation gives room to a number of idiosyncrasies and allows the hearer to manipulate the speaker’s intentions on the basis of what it is convenient or palatable for her to believe. The tension between the individual and the social path of intentionality can be best represented as a tension between selfishness and altruism/responsibility.

In the words of Dascal (2003), communication is regulated by the duty to understand (on the part of the hearer) and the duty to make oneself understood (on the part of the speaker) and the interpretation process is described as a **reaching out towards the speaker**. This terminology is particularly felicitous, because, in describing the speaker’s and the hearer’s duties in the communication process, it emphasizes the notion of responsibility, and the notion of reaching out towards the speaker emphasizes the idea of altruism vs. selfishness as being implied in the notion of a responsible communicator (the social path of intentionality, as Jaszczolt says). As Dascal (2003) says, this reaching toward the other – one might say, this **altruistic orientation** – is inherent in communication qua coordinated action. (Such

considerations are dealt with in detail in four maxims underlying communication in Capone 2004a).

This is part of the picture. The other part of the picture is the attempt to get rid of ambiguity whenever possible. Brandishing Occam's Modified Razor, à la Grice, Jaszczolt proposes that if a lexical item can be interpreted as 'a' and 'b' in different circumstances, instead of positing ambiguity proper, one can say that one is faced with an interpretative ambiguity and one can assign the lexical item a Default Semantics in case one of the two interpretative options is chosen in most contexts. As Mey says:

In real life, that is, among real language users, there is no such thing as ambiguity – excepting certain, rather special occasions, on which one tries to deceive one's partner, or 'keep a door open'. (Mey, 2001: 12).

Take for example the lexical item 'some'. If you utter (1)

(1) Some of the students arrived

the Default interpretation is that 'some but not all of the student arrived' (so some did not arrive, the class is partially empty). If my interpretation of Jaszczolt is correct, her approach is similar but not completely identical to Levinson (2000) on presumptive meanings. And her approach is to be differentiated from the one by Relevance theorists (e.g. Sperber & Wilson 1986; Carston 2002), despite the fact that both approaches can be classified as post-Gricean, intention-based, contextualist accounts of utterance processing. As far as one can see, the main difference between the two enterprises is that while Relevance Theory focuses on processing (see effort/effect), Default Semantics focuses on types of sources from which addressees draw information and types of processes which interact in providing it. It appears that there is scope for an eclectic account. (More on differences later).

Relevance theorists always describe inferences in context. They are mainly concerned with the class of phenomena Grice dubbed 'particularised implicatures'. Default inferences, instead, do not arise from particular contexts – in this sense they are similar to Levinson's presumptive meanings. Jaszczolt's approach is distinct from Levinson's because, while Levinson explicitly ties his inferential

augmentations to pragmatic principles (specifically to his neo-Gricean Horn-based revision of the Gricean maxims), he considers pragmatic augmentations ‘local’. Instead Jaszczolt opts for a global inferential approach, in which inferences are computed by integrating at the utterance level both semantic and pragmatic information.

Jaszczolt gives one the impression that she opts for a semantic view of pragmatic phenomena by proposing the idea of ‘Default interpretations’. As I wrote in Capone (2002), she devises a system in which one stores Default interpretations in an archive and such interpretations are automatically activated in a default context, unless there are visible clues that militate against the Default interpretation and, thus, favor contextual modulation (I do not remember that Jaszczolt makes use of the word ‘archive’ though).

If one wonders how this picture originates, one may probably go back to Wittgenstein’s equation of meaning and use. Alternatively, one may see phenomenology as inspiring Jaszczolt’s perspective – pp. 88-90 of her book ‘Discourse beliefs and intentions’ are instructive in this respect. She takes both Husserl and Brentano to support the idea that the content of our attitudes are things of the world (In Husserl’s work, the connection between perception and belief is emphasized). On p. 48 of her ‘Default Semantics’ Jaszczolt clarifies the issue:

Now, intentional acts can be about mental objects, real objects, or whole states of affairs (eventualities); states, events, or processes. I shall follow the later phenomenological tradition and assume that our mental acts are directed at real rather than mental objects, and at real eventualities. (Jaszczolt 2005, 48).

In this chapter, I would like to give a more distinctively cognitive slant to Jaszczolt’s theory and I would like to claim that there are cognitive principles responsible for the assignment of referential interpretations to NPs. If my view is correct, we need not use premises from phenomenology to support the Default Semantics of NP, nor should we use explanations of Levinsonian or relevance-theoretic inspiration. I propose that the mind works this way. If you encounter an utterance of

(2) Mary thinks that Orcutt is crazy

there are two possible readings:

a. The *de re* reading.

The reporter (of the belief) ascribes to Mary a belief about a particular, known individual (*de re*).

2b. The *de dicto* reading.

The reporter of the belief says that Mary believes in the existence of Orcutt and Mary ascribes to him a certain property.

you will be inclined to give a referential interpretation to ‘Orcutt’ and a ‘*de re*’ reading of the sentence in (2) (Orcutt is such that Mary thinks of him that he is crazy) not because the referential and ‘*de re*’ reading is the most informative one (the one which has greater contextual effects) but because this is the way the mind works. I assume that what Jaszczolt calls the ‘Default *De Re* Principle’ is nothing less than an **a priori form** of our interpretation processes. In the same way in which, for Kant, the ideas of space, time and cause are the *a priori* principles of knowledge, the Default *De Re* principle is an *a priori* form of interpretation processes. The principle is exposed below:

The Default *De Re* principle:

The *de re* reading of sentences ascribing beliefs is the Default reading. Other readings constitute degrees of departure from the Default, arranged on the scale of the strength of intentionality of the corresponding mental state.

- (i) The hearer of an expression of belief of the form ‘B ϕ s’ normally presumes that the speaker holds a belief *de re* and that the referring term is useful to refer to an individual, unless the content of utterance signals otherwise;
- (ii) The hearer of a belief report of the form ‘A believes that ‘B ϕ s’ normally interprets the utterance as *de re*, unless the context signals otherwise.

The Default *de re* principle can be derived through **reasoning**. Suppose that there is no default *de re* principle in language. Then the child, presented with utterances, could very well interpret NPs as not being referential, taking them to refer to abstract categories. For example, faced with an utterance such as ‘The dog is in the garage’ a child could very well take ‘the dog’ to refer to the abstract category ‘dogs’. This means that language would be unlearnable for the child (or learnable with great difficulty), as for every concrete object, the linguistic item referring to it would be multiply ambiguous, capable of referring either to a concrete object or to an abstract category. However, if a child is guided by something like the default ‘de re’ principle, language acquisition is facilitated, as he will not run the risk of misinterpreting concrete categories by applying abstract categories. (Gradually the child will move from the sphere of concrete objects to the sphere of abstract ones, but it is clear that the Default *De Re* Principle is a very useful heuristic principle for the child who needs to orient himself in otherwise chaotic pieces of language behavior).

That children are guided by cognitive principles in language acquisition comes as no surprise, as Carruthers (2006a) claims, following Bloom (2000), that in learning a language children are guided by the innate notion that the speaker’s intention as manifested through gaze or gestures is important; furthermore, he claims that a child will not try to apply a new word to an object part, but to the entire object (So on hearing ‘rabbit’ for the first time, the child will not apply the concept ‘rabbit’ to a rabbit part (say, the hind leg) but to the entire object. (Goldman 2006, 178, discusses the following principle: Prefer parsing the world into whole objects rather than arbitrary parts of whole objects or arbitrary mereological sums of whole objects.)

It goes without saying that the Default *de re* principle is responsible both for the preferred referential interpretation of belief reports and for the preferred referential interpretation of utterances such as:

(3) Smith’s murderer is insane.

Sentences such as (3) are interpretatively ambiguous between a referential and an attributive interpretation: however, the hearer normally settles immediately and automatically for the referential interpretation. Examples of this type can be multiplied *ad libitum*. In fact, an utterance such as (4)

(4) John wants to sell his cello (originally from Heim 1992).

is usually assigned the interpretation that John believes he has a cello and that such a belief is matched by the speaker's belief that John has a cello. On this theory, the case in which crazy John falsely believes he has a cello, even if (4) is uttered, is a marginal one, a reading that constitutes a 'degree of departure from the default'.

Are there any other interesting cases of default inferences? I would like to list a number of inferential phenomena which are not discussed by Jaszczolt, but which may support the case for cognitive Defaults. I will freely resort to cases discussed by Bach (2001) and Dascal (2003), adapting them to the purposes of the present discussion. (I am responsible for modifications)

Speak seriously!

When a speaker A proffers her utterance U, there is a presumption that she seriously intends to say U, **unless** some clues from the context clarify that she cannot have a serious intention (e.g. she is speaking ironically, humorously, etc.). Could we not consider this presumption as part of our cognitive make up – of the way our cognition works? Certainly a cognitive universal to the effect that the speaker's words are taken non-literally or non-seriously unless clues indicate otherwise could not be of any help to mankind; that would only disorient communicators.

The **reasoning** would have to proceed like this. Suppose that there is no default procedure for discerning seriousness of intention. Then the child learning a language would have no guiding principle available orienting her towards the selection of the right lexical items corresponding to appropriate concepts. Language learning would have to be an unachievable task, as for each word, there is the possibility that the

instructor (say, the mother or the father) is not speaking seriously, in which case the motivation for learning a language would decrease.

Only one language!

When a speaker utters more than one utterance, there is a presumption that he would use the same language as that of the previous utterance. You do not struggle to guess, each time: “Which language is she likely to use next?”. In a non-bilingual community, you are pretty sure that a language will be constant throughout the communication process. (In a bilingual community you know in advance that any of the two languages, but NOT a third one, will be selected in actual communication).

The **reasoning** would proceed as follows. Suppose that it is a logical possibility that a sequence of utterances can be uttered by using any language whatsoever, at random. If such a logical possibility were countenanced by the child, the motivation for learning a language would decrease, because, by being exposed to a sequence of utterances, she would not know which language she has to learn and which language the speaker is likely to speak. Instead, if the child expects the speaker to use a single language (or at most a very limited set of alternatives) through a sequence of utterances, the motivation for learning a language would increase.

The speaker as principal

There is a presumption that the speaker is also the **author** and **principal** of his utterance content. These are notions from Goffman’s (1981) *Forms of Talk*. A person can speak without being responsible for the words uttered (e.g. an actor/ an ambassador), he simply voices the message. However, in normal speech we do not dissociate the role of the speaker as the source of the utterance from his role as the means of communication.

The **reasoning** for this cognitive principle would have to proceed as follows. Suppose that there is no such cognitive principle, then for every utterance she hears, the child would not know whether the speaker speaks in her capacity as animator or author or principal. For all the child knows, it is a logical possibility that the speaker is just an animator and that everything she says comes from a source different from

the animator (or the instructor). If such a logical possibility were contemplated seriously by the child, then it would not be possible for her to associate the language spoken by the animator with the principal's language (so the chances that he is learning a language which is not the one spoken by the animator increase), furthermore it would not be possible to associate linguistic items and utterances with the speaker's intentions. But we saw, given what Carruthers says about gaze and gestures, that intentions are important (in fact, indispensable) for learning a language.

These are ideal cases of preferential interpretations expressible through Default Semantics. You do not rely on particular clues – especially words – to draw these inferences. To argue that these are socio-cultural Defaults would imply that these aspects of communication are learnable, rather than innate. But this would require showing that at least in some societies things are not this way. It would also involve explaining how adults/infants communication can occur, since these defaults are **presuppositions** of communication, rather than learnable aspects of communication or of culture. You do not first teach the infant these defaults and then proceed with communication with her; on the contrary, these defaults are presupposed in communication.

We wonder if the cognitive defaults above are reducible to more general principles¹. Now, this question is clearly a question about the link between Default Semantics and Relevance Theory. While a cognitive default may work as an instruction to interpret a certain fragment of language use in a certain way, it is possible that behind it there is a cognitive principle of basic rationality. This I will not deny, albeit I will insist that cognitive defaults are short-circuited inferences, in which the mind is not busy calculating inferences on the basis of general principles of rationality. We can, however, note important connections. Each of such defaults may arise due to the need of avoiding ambiguities and obscurities which would impede not only language processing, but also language acquisition. Since the mind works by promoting contextual effects while keeping efforts as low as possible, and since without such cognitive defaults language acquisition would be impeded or

¹ An interesting reduction of the speaker's 'sincerity stance' to a cognitive story is to be found in Paglieri & Castelfranchi (2010). For the sake of space I cannot go into this discussion.

retarded, the mind recruits Sperber & Wilson's Principle of Relevance for the purpose of creating **cognitive defaults** which, if implemented as simple instructions, are even more frugal and faster than the application of the principle of relevance each time a certain input occurs. We may see the cognitive defaults as **specializations** of the application of the Principle of Relevance.

One may wonder why one should consider these tendencies as best explained as a priori, hardwired principles of (linguistic) interpretation when they can typically be derived from basic rationality. A possible reply to this question is that, granting the connection between these tendencies and the cognitive principle of relevance, which originally motivated them, a person is better off having these tendencies hard-wired, because of the possibility that more than one environmental problem of the type resolvable through them may present itself at the same time. The presence of two or three environmental problems of this type would put inferential reasoning to great costs, since cognition would have to work out responses to two or three or more environmental problems simultaneously. Instead, a hardwired, modular account of these tendencies resolves the problem of simultaneity of inference. This problem, as shown by Carruthers (2006), is at the basis of modular stories. Surely, one could go on arguing that the cognitive principle of relevance can deal with more than one problem in parallel (in other words, it can deal with a series of problems by resolving them in parallel). Now, even assuming that this is so, one must admit that cognitive efforts will increase enormously and there is also a chance that the processor will not be able to cope with so much processing because of the cognitive load. Instead, if we posit that there are some principles related to, but different from, the principle of Relevance that can deal with matters such as 'serious speech' 'author vs. speaker', 'only one language', referential interpretations, the cognitive load for the operation of the Principle of Relevance is diminished, and this is clearly necessitated by the fact that the Principle of relevance is always busy computing the relevance of an utterance in a given context.

It is also natural that the principle of relevance, being a cognitive principle, should be designed to cope with **novel** problems, leaving problems that are **recurrent** to default heuristics, which are prepared, instead, to deal with problems which have repeated themselves in the past and for which 'standardised solutions' are ready. (As Sherry, Schacter (1987) say, it seems reasonable to assume that an efficient habit-learning system will preserve those features of an experience that recur in different

episodes, while a certain number of contextual details are ignored). We can compare the work done by the default semantics heuristics with fast-and-frugal heuristics utilized for example by people who invest money in the stock-exchange. Someone who has witnessed the recent events in the world stock-exchanges has developed a prudential strategy based on experience: if you buy shares when the price fall, buy them in two or three installments, so that if the price continues to fall, you can still be able to buy the shares at the cheapest possible price. If cognition tells you that the best moment to buy shares is when the price falls, experience tells you that if the price continues falling you must go on buying shares. Fast-and-frugal heuristics are the result of cognitive principles being adapted to experiential data. Analogously, I argue that default semantics heuristics are the systematic response to environmental problems that have recurred and, because of their recurring quality, require **dedicated** mechanisms and more specialized solutions. Default interpretations are also more suited to interpretation problems for which there is little doubt or uncertainty, whereas cognitive mechanisms of the inferential type are at work when one needs to eliminate uncertainty. As Gallistel (2002) says, the function of learning is to extract from experience properties of the environment likely to be useful in the determination of future behavior. We use memory to carry information that is useful to responding to environmental problems forward in time.

We may go further and claim that there are other cognitive defaults. Suppose that Default Semantics includes fast and frugal heuristics such as:

What is not said isn't (Levinson 2000).

Given this fast and frugal heuristics, hearing the utterance (5)

(5) I have three children

the hearer assigns the interpretation 'I have at most three children'. This fast and frugal heuristic could be seen as a **specialization** of the principle of Relevance by Sperber & Wilson:

Cognitive principle of Relevance

Human cognition tends to be geared to the maximization of relevance.

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.

Presumption of optimal relevance

- a. The ostensive stimulus is relevant enough to be worth the auditor's processing effort.
- b. It is the most relevant one compatible with the communicator's abilities and preferences.

Relevance-theoretic comprehension procedure

- a. Follow a path of least effort in computing cognitive effects: Test interpretative hypotheses (disambiguations, reference resolutions, implicatures, etc.) in order of accessibility.
- b. Stop when your expectations of relevance are satisfied (or abandoned).

(From Wilson & Sperber 2004).

Some readers may remain perplexed by the fact that I consider the heuristic 'What is not said isn't' as related to the Principle of Relevance, nevertheless as belonging to Default Semantics. However, given Carruthers's claim that different modules can share a number of parts, I also claim that Default inferences and the Principle of Relevance share common mechanisms (the relationship between the two being one of specialization). The heuristic principle 'What is not said isn't' can be said to follow from the principle of Relevance, given that if the scalar interpretative options higher up in the scale obtained, then the speaker would put the hearer to undue processing efforts, in that one could never know in principle which of the two interpretation options should be chosen (hence the higher options in a scale are

excluded). As Mey (2001, 69) says: when communicating, speakers try to be understood correctly, and avoid giving false impressions.

I have reasons for wanting to propose that the heuristic ‘What is not said isn’t’ is part of Default Semantics: 1) it furnishes pre-contextual processing; 2) it is a more frugal principle than the principle of relevance.

3. On the compatibility of Relevance theory, Gricean theories, and Default Semantics

The project of integrating Gricean theories and, in particular, Relevance Theory and Default Semantics is not without problems. There are, in fact, terminological, as well as more substantial differences. If Grice’s original theory is taken into account, the most obvious differences between the original Gricean picture and Relevance Theory are as follows.

The Gricean picture is philosophically motivated, while RT aims at a psychologically plausible theory in which Relevance is part of a broad picture of the way the mind works and of its cognitive architecture. Grice had a view of implicatures that had literal meanings as a point of departure, while RT abandons the idea that the literal content gives rise to implicatures, which may be aborted in context. Grice had a minimal conception of what is said, while RTs have an intrusionistic view of what is said. RT adopts the idea that semantics underdetermines pragmatic interpretation to a greater extent. (See Horn 2005, Carston 2005). *En passant*, I should stress that despite much insistence on the part of the current literature on the different emphasis of the Gricean project and RT concerning the speaker’s intentions, Carston (2005) has (as always) said that RT is very much concerned with speaker meaning, both what it is and how an addressee attempts to recover it. (And I approve of Carston’s precis).

We may agree with all this, and say that while we are obviously interested in the Gricean notion of speaker’s meaning, which we hope to inherit for the purpose of our current intellectual enterprise, we are less interested in rehabilitating the less

tractable parts of the Gricean program. Instead, I am interested in integrating, if possible, an offshoot of Gricean pragmatics – Default Semantics – with Relevance Theory. Here too compatibility is not to be taken for granted, as there are differences. However, one should not think that these differences will prevent us from contemplating the possibility of a unified picture. In the following I do my best to reconcile the two positions. I focus on the differences between the two projects and argue that these differences cannot prevent us from trying to reconcile the two theories.

While Relevance Theory focuses on processing (see effort and effect interplay), Default Semantics focuses on types of sources from which addressees draw information and types of processes that interact in providing it. In Default Semantics, rationality of interlocutors is assumed, on a par with any Gricean approach and Relevance Theory. Rational communicative behavior is an assumption integral to the theory of Default Semantics.

Like Relevance Theory, there is a difference between DS and Levinson's theory of GCI in that Levinson assumes local inferences, often word-based ones, while DS professes 'methodological globalism': it is methodologically prudent to assume global, proposition-based inferences (and defaults).

Inferences, according to DS, are NOT fast and instantaneous. DS is a theory in which merger representations combine outputs of various processes, some of them default and some conscious, costly and inferential. Duration of processing is not a problem for DS: GCI will arise in some contexts as defaults, in others as the outcome of conscious processing, and yet in others do not arise at all. Defaults mean in DS automatic interpretations, assessed for a particular utterance. We must distinguish between Default Semantics as a theory and default interpretations (CD, SCWD) which constitute only two out of four types of processes in the theory. DS is a model of utterance interpretation which is self-contained and comprises also those types of processes that Relevance Theory postulates (contextual inferences).

Relevance theory is clearly and explicitly related to mental heuristics. The framework deals with cognitive processes that tend to reduce the amount of information needed and produce effects in a fast and efficient way. If you asked me

whether Miland is bigger than Caltanissetta, I may use information that Miland has a series A soccer team while Caltanissetta hasn't as a predictor of size and if I do not have information about soccer then I will move to the next best predictor of size (for example, having an airport is a predictor of size). In the same way, Relevance Theory uses cognitive efforts and contextual effects to predict whether an interpretation is relevant or not. However, one could say that while the heuristics such as using a predictor of size (e.g. knowledge about a soccer team) are intelligent, they are nevertheless shortcuts grounded in experience, Relevance Theory interpretations are shortcuts grounded in general principles of cognition. If anything at all resembles Gigerenzer's heuristic shortcuts, this is Default Semantics and not relevance theory, since Default interpretations are often derived by default, and often include social defaults. They are shortcuts that are grounded on experience and associative links, given that sociocultural defaults are created by associative links.

Default Semantics must be related to cognitive architecture because, since it provides principles of compositionality for the integration of linguistic semantics and pragmatic information, one must explain what features of cognitive architecture permit inter-modular communication. As previously said, DS is a theory in which merger representations combine outputs of various processes. It is Default Semantics that permits the integration of, say, perceptual and linguistic information or of syntactico/semantic and pragmatic information. While Relevance theory confines itself to adding pragmatic increments to semantic templates, Default Semantics gives us the principles whereby such integration can be effected. Carruthers (2006) discusses how information coming from different modules can be integrated, and mentioning that some scholars prefer the answer that the theory of mind is responsible for such integration, he proposes that the language module is responsible for this integration. If we assume, as Jaszczolt does, that Default Semantics is responsible for this integration, we must surely articulate mental architecture that corresponds to the workings of Default Semantics. Since Default Semantics is said to be characterized by combinatorial principles, it is not completely outlandish that the workings of Default Semantics should subsume a linguistic competence, but it is also obvious that such combinatorial principles must be different from linguistic compositionality.

What is clear so far is that an obvious difference between default Semantics and Relevance theory is the former's insistence on Default inferences. These include socio-cultural defaults, which cannot be easily accommodated by Relevance Theory. A clear-cut advantage of Default Semantics is that it can explain inference to stereotype (nurse → female nurse; surgeon → male surgeon, etc.) with relative ease. Default inferences of a lexical type are also a characteristic of Default Semantics, which is in conflict with the assumptions of Relevance Theory. In my opinion, one must attempt to integrate DS and RT by recognizing that much of the burden of cognitive processing must be alleviated by the construction of a mental archive where Default inferences are stored, ready for use. Even if it is not clear that this approach involves faster mechanisms of inference, one thing is certain. Processing costs will be minimized, and this is useful especially when the brain is at pains locating an utterance in a wide co-text and context. If co-text and context are not at odds with Default interpretations, these can be safely added to utterance interpretation by following the Default Semantics compositional principles and can be integrated in a wider context where it is necessary to compute further contextual effects (the effect of adding the utterance to the context in question). Computing costs will thus be predicted to be alleviated by Default Interpretations. These considerations make sense in the light of Perkins's (2007) work on compensatory strategies, which clearly involve brain plasticity as well as the natural inclination by the human mind to compensate for certain impairments. While in the context of this discussion we are not discussing pragmatic impairments, it is natural to extend the notion of compensation to those cases in which the summation of inferential procedures would ultimately lead to an unbearable or anyway costly cognitive load and, thus, brain plasticity, and, in particular, the ability to shift information from the cognitive component to a Default Semantics Archive, takes away this cognitive load and compensates for it by imposing a burden on the Default Semantics Archive.

In order to be more persuasive, I will make use of an example which is familiar to everyone: the use of the multiplication tables. I choose this example from Ryle (1949), since the philosopher used it to make a distinction between an intelligent activity (multiplying numbers according to a rule) and an activity which seemed to him to be less intelligent (the ability to give by rote the correct solution to a multiplication problem). We do not know what Ryle would have said of the strategy of shifting the burden of cognitive operations to a semantic archive in which semantic associations replace cognitive operations. Would this be an intelligent

activity? If we could show that this activity is governed by the principle of relevance, then we could prove that it is intelligent. I will address this issue later on. When we were children we found it tedious to commit to memory so much information about multiplication, which was independently derivable through a cognitive strategy (e.g. counting, or multiplying with the usual multiplication methods). This clearly involved a cognitive load on memory. However, once the multiplication tables were memorized, it was clear that a lot of effort was saved. Of course, committing information to memory is useful only if there is a high likelihood that this information will turn out to be useful. Thus, if in our practice (buying and selling things, measuring, etc.) there is a high likelihood that the multiplication tables will turn out to be useful, it will be felt to be justified to put children to such a memory load. Could we not then arbitrarily extend the capacity of multiplication tables? I found out through Wikipedia that while in some society multiplication tables include the number 9, in USA multiplication tables include the number 12. However, one usually does not find multiplication tables that include the number 19 or 99. Why should this be the case? It must be the case because given that it is unlikely to find multiplications like 99×88 in practical life, then the cost of memorizing such richer multiplication tables is not justified. Even for mathematicians, this is not justified, given the use of electronic calculators. However, if a given person found it was very useful to multiply numbers up to 99 in practical life (say in a society of thieves where everyone will take an opportunity to steal money from you if you do not have suitable mathematical abilities), then it would be justified to learn the multiplication tables up to the number 99.

When we have to choose between a simple associative cognitive strategy, like memorizing, and a more complex cognitive strategy, such as having to make computations, we compare the costs and benefits involved and then decide which strategy is the best. We can invoke the basic assumption of relevance theory, that is the trade-off between cognitive effects and cognitive efforts to justify the shift from repetitive processing within the module where the principle of relevance is operative to systematic storage in a Default Semantics archive of information coming from processing based on the Principle of Relevance. Furthermore, if ease of memorization is associated with repetition, then it goes without saying that the more repetitive a cognitive strategy is, the greater is the likelihood that it will be replaced by an associative learning and memorization of its results. Suppose there is a cognitive rule of this type: after finding that item x is associated with interpretation y

for the n th time, commit y to memory by associating y with item x . Then it goes without saying that after encountering the association x, y for the n th time, one will commit it to memory. However, at this point we could also say that Relevance Theory mechanisms have some saying on the fact that the rule ‘Commit the association z, y to memory after encountering it for the n -th time’ exists, since memorization effort is now offset by cognitive rewards. (This is in line with Millikan 2006, who also stresses the connection between memorization of facts and cognitive rewards). All I am suggesting is that the number n involved in the expression ‘the n -th time’, whatever it is, is not arbitrarily determined, but may be the result of a combination of factors, including the dispositional qualities of memory (synaptic density in neo-cortical tissues, changes in synaptic conductance (Gallistel 2002)) and cognitive rewards that are likely to offset memorization effort.

There is more evidence in favor of the line of thought according to which shifting the burden from on the fly cognitive processing to memorized information (the Default semantics archive) is liable of functional explanation. I would like to suggest an analogy between the Default semantics Archive /on the fly processing tension and the functional considerations by Klein et al. (2002) on episodic memory and ascription of character traits. Essentially Klein et al. consider that ascription of character traits can be determined in two ways: either through online processing, in which one considers episodic memories of the past and ascribes qualities on their basis; or through transformed memories, in which memories are elaborated on and abstractions are made on their basis. Such abstractions provide character traits ready for use. The preference for character abstractions is due to the fact that they respond efficiently and quickly to situations in which a decision and an action must be taken in a fraction of a second. The reasons for this cognitive default are mainly the following: predictability (one can predict another person’s future action); importance (the importance of decisions should justify the costs); economy (the number of judgments supported by derived memories should be proportionately large with respect to its size; urgency (the requirement that such judgments should be made quickly).

There is a tension between abstractions based on memory of events and online processing based on events, since abstractions have predictive power concerning average situations. However, when the question arises how a person should behave

in an exceptional situation, it is of use to have a memory store that keeps memory of events, particular episodes, to see if such a memory can tell us what to predict about future behavior.

The tension between online processing of memories and memory based abstractions of character properties parallels the tension between online processing in which a relevance based mechanism is able to compute the effects of context on an utterance and the information stored in the default semantics Archive. Default inferences are of use when cognitive costs increase due to parallel processing of various lexemes and structural configurations. The efficiency of Default Inferences should, therefore, **justify** the costs. Default Inferences have **predictive** power. They tell us how subjects will behave conversationally and use lexical items standardly. Default inferences involve a certain **economy**, because it is presumed that an inferential augmentation is routinely used million of times in a person's life and, thus, cognitive efforts involved in producing the inference are spared; default inferences also meet Klein et al.'s criterion of **urgency**. If default inferences are involved in thought as well as in spoken language, then it is clear that they allow speakers/thinkers to act quickly.

To finish this defense of Default inferences, I want to say that when a cognitive resource is very fruitful, as in the case of pragmatic inferences, then it will be good to make sure that other cognitive resources are allocated to the same function, so that this function will not be lost as a consequence of neurological damage. In the same way as our perceptual systems have got a certain redundancy (we have got two hands, two eyes, two ears, a nose and a mouth that can both be used for breathing, two legs, two feet), our cognitive resources are best used by creating redundancy and by replicating them by shifting away the burden from cognition to memory.

A further crucial difference between DS and RT is that unlike the other contextualist account, DS does not recognize the level of meaning at which the logical form is pragmatically developed/modulated as a real, interesting, and cognitively justified construct. To do so would be to assume that syntax plays a privileged role among various carriers of information and that the syntax/pragmatics interaction is

confined to pragmatic additions, embellishments, or ‘developments’ of the output of syntactic processing.

Now, it is possible that this insistence on the idea that semantics is not a cognitively justified construct needs to be mitigated. After all, the idea of merger representations is that we merge information and linguistic information is clearly one of the sources of information that we merge. What Jaszczolt probably wants to challenge is the independence of semantics. Now, there are reasons for resisting the idea that semantics is an autonomous level of meaning, since semantic representations, in order to be merged, must have tags that allow for the merging. In other words, linguistic information must have appropriate tags that allow the binding (or merging) with non-linguistic information (The simplest example is the use of pronouns or demonstratives). In other words, they must have a potential for merging. Seeing things in this light, the insistence by Relevance Theorists on semantics as an autonomous, independent, cognitively justified level can be tempered. However, given that Jaszczolt deals with compositionality at the level of the utterance, we should remark in passing that having the potential for compositionality at the level of merger representations does not prevent semantics from being **readable** as a self-sufficient level of representation. One knows what could be meant by saying a given sentence and one also knows in advance what increments various and different contexts would add because we know what the semantic potential for merger representations is. It is probably this abstract semantic potential for combining with contexts that we call sentential semantics and there is no doubt that Relevance Theorists do well in paying attention to this level which may require different principles of compositionality from those of Jaszczolt’s merger representations

3.1 Default Semantics and Levinson’s theory of default implicatures.

While there is no doubt that Jaszczolt’s theory of Default semantics seems (prima facie) to resemble Levinson’s (2000) theory of generalized conversational implicatures, in the light of evidence from experimental pragmatics (especially Noveck & Sperber 2007), we should make some effort to differentiate the two theories. Noveck and Sperber (2007) note that default implicatures à la Levinson presuppose the following theoretical picture:

Generalized implicatures are more parsimonious than totally explicit communication;

Generalized implicatures are more parsimonious than particularized implicatures (in the sense that they are less time-consuming, less effortful).

Generalized implicatures get through in a default context, but can be cancelled in specific contexts;

The contexts where they are cancelled are fewer than those in which they are not.

Noveck and Sperber (2007) consider that even if generalized implicatures were cancelled only in a third of all cases, one would have to sum the effortfulness of these cancellations with the effortfulness of having generalized implicatures. The criticism leveled against Levinson could be leveled towards Default Semantics. Thus, we must excavate how to answer the objections to Levinson while at the same time differentiating Default Semantics further from the Levinsonian picture.

3.2 Taking experimental pragmatics seriously.

The experimental side of Noveck and Sperber's article seems to show 'prima facie' that children are more likely to go for literal interpretations of utterances such as 'There might be a parrot in the box' (which are preferred to interpretations like 'There might but there must not be a parrot in the box' (See also Chierchia et al. 2001 on this issue). The experiment is not much different in its results when children are given explicit instruction concerning the presence of a scalar conversational implicature. Now, Noveck and Sperber agree that these results can be interpreted by saying that all they show is that pragmatic competence in children and in adults are differentiated (does the theory of mind module undergo a paced pattern of evolution?) and, in fact, the temptation to interpret them as giving support at least to a picture of Default Semantics à la Jaszczolt is strong. If my idea is correct that information is gradually shifted to the Default Semantics Archive, after inferences become routinized or standardized, say after the *n*th time they are repeated, then the experimental picture so far illustrated need not be inimical to Default Semantics, as the Default Semantics Archive could be considered a component of the theory of mind which partially **emerges** with experience. However, the literature on experimental pragmatics is not unanimous on the idea that children are not pragmatically competent to calculate scalar implicatures. In fact, Papafragou and Tantalou (2004) argue that the methodology of the previous experiments, in which questions were asked of the children concerning the truth of certain statements, was wrong and, thus, use instead a different methodology which does not privilege

questions, but exploits informative statements containing quantifiers to assess whether a certain task was completely executed and thus deserved a reward (or not). In the light of the new methodology, children were shown to be competent in calculating scalar implicatures **if** enough contextual clues were given.

Another important article belonging to experimental pragmatics is the one by Breheny et al. (2006) in which it is shown that contextual clues may lead either to scalar implicatures or to a literal interpretation respectively in greater or less time. Now, this is certainly an important study of inferential processes, but there are things to be said. If understanding a sentence is a question of connecting together its various parts, it is clear that syntactic complexity should be taken into account. It could be argued that syntactic complexity of prior segments have an effect on the reading time of the last segment, given that, after all, reading in this experiment meant understanding and understanding is always a holistic matter.

The problems noted above do not appear to frighten Jaszczolt who distances herself from Levinson and claims that:

It is much harder to provide experimental evidence for or against salient meanings that are so construed that they draw on some contextual information, arise late in utterance processing, and are not normally cancellable. The latter also seem much more intuitively plausible in that they are nothing less but shortcuts through costly pragmatic inference. They are just normal, unmarked meanings for the context and it is not improbable that such default, salient interpretations will prove to constitute just the polar end of a scale of degrees of inference rather than have qualitatively different properties from non-default, clearly inference-based interpretations. They will occupy the area towards the 'zero' end of the scale of inference but will not trigger the dichotomy 'default vs. inferential interpretation'. (Jaszczolt 2006).

3.3 Further considerations on the interaction between Relevance Theory and Default Semantics.

One of the ideas of this chapter is that inferences that start as pragmatic processes become standardized (to use Bach's term) or routinized. This involves a gradual shift of information from pragmatic processes based on fast and frugal heuristics to a (lexical) archive. This move is not extraneous to Relevance Theory as evinced from Wilson and Carston's (2007) brilliant paper. In that paper, Wilson and Carston attempt to unify diverse pragmatic processes such as loosening or narrowing of concepts, or metaphorical extensions, hyperbole, etc. They convincingly argue that there is no clear division among them. They can successfully unify such diverse phenomena under the idea that Relevance will furnish 'ad hoc' concepts (an idea based on Barsalou 1987 and Glucksberg 2003). In other words, pragmatics will be able to modulate in context the concepts provided by lexical semantics. In that paper, they also argue that "However, some of these pragmatically constructed senses may catch on in the communicative interactions of a few people or a group, and so become regularly and frequently used. In such cases, the pragmatic process of concept construction becomes progressively more routinized, and may ultimately spread through a speech community and stabilize an extra lexical sense". (Wilson & Carston 2007, 15). As Mey says:

(...) certain apt metaphors (e.g. 'sharp' for 'intelligent'), due to their 'success', obtain near-lexical status, analogous to certain fixed expressions (compare the role of the English modal verbs *can* and *may* in indirect speech acts and negation) (Mey 2004, 113).

Now, it is clear that both Default Semantics and Relevance Theory recognize the importance of routinized inferences (see also Mey 2004, in addition). One may, thus, object to Default Semantics that, after all, if all it does is to create an archive for the new senses that seep into the language, it amounts to no more than something like the lexicon as classically conceived of. One can reply to this by saying that, of course, Default Semantics is more than a lexicon. I have already said before that it contains some very useful Cognitive Defaults, in additions to principles for the

calculation of merger representations. Relevance Theory provides reasoning principles and Default Semantics the mechanism and the algorithm where compositionality applies to utterances being a product of interaction of information from various sources. So Default Semantics is a set of cognitive defaults and principles determining merger representations, but also has an archive which is very specific. While the lexicon, as normally intended, contains lexical rules that support monotonic inferences (If every student passed the exam, then some student passed the exam), the Default Semantics Archive contains Default interpretations, interpretations which can be defeated in a certain context.

Although Default Semantics does not specifically address the issue of metaphors, it is clear that this is the area where the interaction between Default Semantics and the Principle of Relevance is more evident. Consider the cases of incandescent metaphors used to speak about Berlusconi in the recent political debate. Surely, people who utter the following literal translations of Italian sentences ‘one must get rid of Berlusconi’, ‘one must give a punch to Berlusconi’, “one must stop Berlusconi”, etc. are not speaking loosely, but metaphorically. Such metaphors, especially in the Italian translations of these sentences, have become so standardized, that one would hardly say that they are incitements to violence, as the right-winged politicians immediately argued after the recent physical attack on Berlusconi. However, it is also true that in a discourse where many of these metaphors are used, the conventional metaphors may be transformed into ‘ad hoc’ concepts again through Relevance heuristics (a deactivated concept can be reactivated).

It is a pity that Jaszczolt does not say much about metaphors, but in principle it is possible to shift under the rubric of Default Semantics many of the data on metaphors by Giora (2003). The main idea is that (conventional) metaphors do not allow the hearer’s access to the literal meaning first, but they are cases in which literal interpretations take more time than non-literal interpretations. They are, therefore, ideal cases for Default Semantics.

To defend the idea that conventional metaphors are stored within the Default Semantics Archive one needs to argue the case that conventionalized metaphors are mid way between genuinely pragmatic processes and lexical inferences. A recent debate on the language used by the media and politicians (prior to a violent attack on him) in which Berlusconi was described by (violent) metaphors and the charge that such a language was intentionally used to create a climate of violent opposition, proves that the idea of storing conventionalized metaphors in the Default Semantics

archive is not outlandish, since these metaphors still live in some intermediate stadium between what counts as *langue* and language use.

4 Modularity and innateness

I briefly discuss the question whether the Default Semantics Archive is completely innate or, otherwise, developed due to exposure to experience. My answer is of a mixed type. On the one hand, heuristics such as the Default De Re principle seem to be **a priori principles**, hence innately present in the mind, or at least related to the cognitive principle of relevance – which is an a priori principle. On the other hand, the Default semantics Archive grows up as a result of exposure to experience. It is similar to the mental lexicon, in the sense that experience leads to progressive accumulation of information – of course, some ordering principles must be present in the archive, and these are presumably innate. The growing of the Default Semantics archive resembles very closely the modularization process involved in skills such as driving and reading. According to Barrett & Kurzban (2005) and Karmiloff-Smith (1992) the modularization involved in novel skills such as reading recruits evolved modular capacities such as object recognition. The experience of reading would influence the development of the reading-module in such a way that the developed system, as observed in reading adults, appears to contain a specialization for reading. According to Barrett & Kurzban, it is possible that while the reading module recruited the object recognition module to start with, there was a **bifurcation** of modular skills during development. As Barrett & Kurzban say:

In this case novel tasks such as identifying letters or words, would still be treated by the evolved developmental system as a special case (or token) of an evolved skill (object recognition) if they satisfied its input criteria. However, the development system in question could contain a procedure or mechanism that **partitioned off** certain tasks – shunting them to into a dedicated developmental pathway – under certain conditions, for example when the cue structure of **repeated instances**² of the task clustered tightly together, and when it was encountered repeatedly, as when highly practiced (...). (Barrett & Kurzban 2006, 639). (Bold mine)

² Bold mine, in this case.

The considerations above seem to me to be of extraordinary importance. The Default Semantics archive could be seen as originally recruiting information from the operation of the Principle of Relevance, **until** the tasks it carried out started to be partitioned from the mechanism of Relevance in such a way as to allow for dedicated mechanisms and specializations causing ease of cognitive load. If the ontogenesis of the Default Semantics archive reflects that of the reading module, we expect dissociations between the mechanism of Relevance and the Default Semantics archive as a result of the bifurcation in cognitive specialization that resulted in a new sub-module. In other words, my expectation is that damage to the Default Semantics Archive may leave the operation of the Principle of Relevance untouched and vice versa. This dissociation is, of course, ultimately very valuable for the system: suppose that the Default Semantics Archive, for some reason, ceases to exist or to function properly; then its inferential power will have to be replaced by the less dedicated Relevance heuristics.

The considerations by Barretht & Kurzban also support the existence of the Default Semantics Archive as a separate sub-module in that the accumulation of Default Semantics information in the dedicated archive is the **result of repeated practice**. We may suppose that there is a mechanism specifying that, if a certain inferential procedure is routinely used, and its repeated use distracts cognitive resources from the computation of relevance, then the output will be systematically stored in the Default Semantics archive exploiting the Principle of Relevance.

Conclusion

The most important idea of the chapter is to link the interaction between Default Semantics and the Principle of relevance to a modular picture of the mind. In fact, I have sketched with broad brush strokes a picture that integrates Relevance Theory and Default Semantics within the same modular architecture (the mind-reading module) and I have advanced the hypothesis that the Default Semantics heuristics share mechanisms with the Principle of Relevance, as, in fact, they originated ontogenetically by originally recruiting cognitive operations from that mechanism, which, once routinized, would turn into cognitive defaults which would then split off from the mechanism of relevance (the two mechanisms, that originally shared resources, subsequently **partitioned** into two separate and dissociable mechanisms.

Chapter 3

What can modularity of mind tell us about the semantics/pragmatics debate?.

Abstract

This chapter is mainly about cancellability of explicatures and modularity of mind. Cancellability is the property of inferences whereby a certain inferential output, when confronted with further evidence coming from a different archive, is aborted. World knowledge, in particular, has the ability to suppress a certain inference. Cancellability is taken by Kasher as conflicting with the idea of modularity and showing that pragmatic processes (of an inferential, defeasible type) cannot be properly called modular. Kasher's view must be reassessed on the basis of my ideas on non-cancellable explicatures, expressed in Capone (2006, 2009a). The idea that explicatures cannot be cancelled ties in very well with the idea that heuristic strategies, such as those which guide inferential processes, are satisficing strategies.

Introduction

In this chapter I shall make connections between two domains of information, research on the semantics/pragmatics debate and on modularity of mind, in the hope that establishing connections and parallel structure may be fruitful in deepening knowledge of the interface between semantics and pragmatics. In particular I want to inquire if modularity of mind can help us move towards the resolution of important theoretical problems like Grice's circle, the cancellability of explicatures/implicatures, the analogy between perceptual enrichments and explicatures due to free enrichments, the routing problem for explicatures (do they strictly take input from implicatures?), the distinction between primary and secondary pragmatic processes. I hope that something new will emerge out of this system of parallel inquiry, that will, if not provide definitive answers to our search for the truth, at least lead us to more solid landmarks, from where to advance further.

Before moving on to the core issue of the chapter, I shall make a detour and expatiate on the semantics/pragmatics debate and on modularity of mind. In proceeding so, I shall give ample space to my papers on Grice's circle and on the cancellability of explicatures (Capone 2003, 2006, 2009a), as my main hope is that this chapter will be able to provide further corroboration of those ideas – or at least to advance that discussion in the right direction.

One of the tenets of modularity is that the operations of a module are mandatory and encapsulated. Thus, if inferential processes are modeled on the basis of modular processes, they are also mandatory and encapsulated. This means that they are triggered automatically by an input and that they cannot have access to the operations of other modules. These two properties have been taken by Kasher (1991) to imply that pragmatic inferences should NOT be cancellable. Cancellability, in fact, is the property of inferences whereby a certain inferential output, when confronted with further evidence coming from a different archive, is aborted. World knowledge, in particular, has the ability to suppress a certain inference. Cancellability is taken by Kasher as conflicting with the idea of modularity and showing that pragmatic processes (of an inferential, defeasible type) cannot be properly called modular. Kasher's view must be reassessed on the basis of my ideas on non-cancellable explicatures, expressed in Capone (2006, 2009a). The idea that explicatures cannot be cancelled ties in very well with the idea that heuristic strategies, such as those which guide inferential processes, are satisficing strategies. When one finds an optimal solution, one stops the inferential process and does not start it again in the light of further inference. While inferential enrichments are open in principle to information coming from a variety of knowledge archives, the Principle of Relevance bounds the search for information creating a smaller module (I would say, a module **on the fly**¹) from which interpretative options are chosen and where completing information for the saturation of incomplete radicals is drawn. For example, in the case of 'John is ready' the search for information from which the

¹A **module on the fly** is a module that limits the information that can be searched in order to arrive at a certain interpretation. For example, consider the case "The surgeon was ready to operate John. And he was ready to be operated". The module on the fly posited by Sperber & Wilson restricts the database where referents can be searched to the entities mentioned in the module [The surgeon was ready to operate John] – as conversation proceeds further, further modules are built on the fly, which are very local in nature.

incomplete radical requires completion is information attending to the main salient goal of the discussion – so if the conversationalists are discussing an exam, or the goal ‘John wants to take and pass the exam’ the search for missing information is confined to an on the fly module that takes into account at least a goal partitioning a slice of relevant information from the rich structure of the previous conversation). This make the inferential process and instantaneous one. A good idea is to distinguish between instantaneous inferences, which are part of modular processes, and non-instantaneous inferences, which are modeled after explicit and conscious reasoning.

Another idea of the chapter is that pragmatic processes often consist of top-down effects, which need to be combined with bottom-up effects. This idea seems to tie in very closely with the enzymatic theory of modularity, which abandons the idea that inter-modular connections are pipes, allowing only a one way traffic, but voices the idea that modules are linked by enzymatic processes which communicate with a common bulletin board. This common bulletin board contains both inputs and outputs. The output of a process then becomes the output of another process. Enzymatic processes are processes in which the formal properties of an input determine the binding of the input to an enzyme-like element that makes the input available for transformation and then delivers it transformed, as output, to the bulletin board. Completion processes involved in the semantic/pragmatics debate are clearly bottom-up processes which must interact with top-down processes. A model based on a one way traffic, through pipes, cannot do justice to this complex interaction. Instead, a model based on a common bulletin-board can show how completion processes and free enrichment processes (which are characteristically top-down) can interact.

1. The semantics/pragmatics debate

We may write many pages on the semantics/pragmatics debate, but I will confine myself to four positions:

Relevance theory (as exemplified by Carston);

Moderate Contextualists (as exemplified by Stanley);

Default semantics, as exemplified by Jaszczolt. (For a more comprehensive view of the semantics/pragmatics debate see Capone 2006, 2008, 2009a).

The considerations by Capone on intentionality and cancellability of explicatures.

Carston (1999) develops the views of Relevance Theorists such as Sperber & Wilson (1986) and considers cases of free enrichment where the logical form provided by the semantics of the sentence is enriched (freely, without the aid of empty categories projected by the syntax or by the semantics)) through pragmatics (also see Hall 2008). Such enrichments result in the ability to provide the utterance with full truth-conditional specifications (it is claimed that in many cases without such inferential enrichments a fragmentary or patently false proposition is provided by sentential semantics). The following cases exemplify free enrichment:

Completion cases

- (1) Paracetamol is better [than what?]
- (2) It's the same [as what?]
- (3) He is too young [for what?]
- (4) She is leaving [from where]
- (5) It's raining [where?]

Expansion/strengthening cases

- (6) Everyone went to the party;
- (7) He took off his boots and got into bed;
- (8) Writing an essay will take time;
- (9) He hasn't had lunch.

The completion cases above are resolved by answering the Wh-question in parentheses. The expansion cases are resolved by expanding the logical form and providing an element which bounds its truth conditions providing reasonable interpretations. Thus in (6) the speaker confines himself to people from a certain domain (e.g. colleagues); in (7) temporal relations between the events are fully specified; in (8) the trivial literal proposition is turned into a less trivial proposition by specifying the lapse of time required; in (9) the obviously false proposition is

turned into a possibly true proposition by providing a reasonable specification of the time in relation to the question (9) is a reply to.

The position by moderate contextualists like Stanley (2007) uses the assumption that pragmatic interpretations are constrained by elements of the logical forms of sentences. Such logical forms posit variables, which are seen as part of syntactic structure (empty categories, in other words). I provide below a typical example of a semantic defence of hidden indexicals by Stanley.

According to Stanley, a sentence such as:

(10) Every bottle is in the fridge

corresponds to the structure:

(11) Every [bottle $f(i)$] is in the fridge

where $f(i)$ is a function from a domain of quantification to objects in that domain. In context, one specifies specific values of $f(i)$ – say the set of bottles I bought at the supermarket, etc.

Stanley supports his analysis through syntactic evidence, most notably from binding:

(12) In every room, every bottle is in the corner

has a structure corresponding to (13)

(13) In every room r , every bottle in r is in the corner.

However, as pointed out by Carston and others, this argument over-generates, as one could argue that in the sentence (12) there is also a variable standing for a domain:

(14) Everywhere I go, $2 + 2 = 4$ at x

Stanley defends his position, by saying that, since it is obvious that there is not a variable for a domain in (14) one should not posit one.

Concerning Stanley's specific analyses, I will merely say that the general idea that semantics constrains pragmatic interpretation can be safely adopted, and can be reconciled with the spirit of Relevance Theory, which is often called a dual pragmatic theory because of its lavish concessions to semantics (Borg 2004).

Jaszczolt's theory is far more radical. Jaszczolt (2005) posits that merger representations integrate semantic and pragmatic information, including socio-cultural defaults. Jaszczolt's merger representations are not obtained by an orderly inferential incrementation, implicatures being able to shape explicatures and explicatures being able to shape implicatures. Jaszczolt's merger representations include both bottom-up and top-down effects. Her merger representations constitute the level where the principle of compositionality is met, given that such a principle is violated by semantic representations (a typical case is instantiated by belief reports (see Jaszczolt 1999; see also Capone 2006, Capone 2008a) This is not the place for a deep discussion of these ideas, but I want to stress that the idea of merger representations ties in very closely with a certain perspective on the modularity of the mind.

2. On cancellability of explicatures

In Capone (2006) and Capone (2009a) I discussed various cases of pragmatic intrusions. Pragmatic intrusions are inferences relying on the same processes on which conversational implicatures are based and contribute to the proposition expressed. (Without a pragmatic process, there would be no proposition to evaluate, but only a fragmentary sentential schema). They are usually associated with a certain construction, called 'intrusive construction'. For a deep discussion of this issue see Huang (2007). The cases of pragmatic intrusion listed there are: reference resolution, deixis fixing, disambiguation, ellipsis unpacking, generality narrowing, comparatives, disjunctions, because-clauses, but it is clear that the forthcoming literature is eroding more and more the terrain of semantics proper in favour of truth-conditional pragmatics. In Capone (2009a) I reflected on the necessity of pragmatic developments of propositional forms in order to arrive at a better understanding of

the level of meaning which Sperber & Wilson (1986) and Carston (2002) call ‘explicature’. In this paper I argued that the pragmatically conveyed elements of explicatures are not cancellable – unlike conversational implicatures. While Capone (2003, 2006) addressed the issue of the cancellability of explicatures from a more empirical point of view, in this chapter a number of important theoretical questions are raised and discussed. In particular, it is proposed that the analysis of the notion of intentionality and the nature of pragmatic intrusion will settle the question of the cancellability of explicatures. An explicature can be considered a two-level entity. It consists of a logical form and a pragmatic increment that the logical form gives rise to in the context of utterance. However, both the initial logical form and the pragmatic increment are the target of pragmatic processes. Consequently, we need a pragmatic process to promote the initial logical form to an intended interpretation and another pragmatic process to derive further increments starting from the initial logical form as promoted to an utterance interpretation.

So in Capone (2009a) the main claims rest on 1) intentionality) the nature of pragmatic intrusion. In this resume I will confine myself to these two claims.

2.1 Intentionality.

The reason why particularized implicatures cannot be cancelled is the same that prevents explicatures from being cancelled. When a strong intentionality is projected, it can no longer be retracted². Implicature can only arise if intended and recognized as intended. But then it should be impossible to cancel an implicature: how would it be possible to withdraw/cancel what was intended to be implicated and was recognized as intended? An implicature could only be withdrawn/cancelled if it were NOT intended. But then it shouldn’t BE an implicature (since implicatures by definition are intended); in other words there would BE no implicature to cancel. This is relevant to the issue of explicature cancellation only in this: presumably, explicatures must be intended and must be recognized as intended. So the above

² An interesting potential objection is the following. In the case of Grice’s well-known example (A: Has John got a girlfriend? B: He has travelled to London very frequently recently) the implicature that John has got a girlfriend in London can be cancelled by adding ‘I don’t intend to mean that he has got a girlfriend in London’. This is surely a case of weak commitment of an implicature (not to mention the vagueness of the implicature; a lot more could be implied as well, and it is NOT clear here that a certain intention can be pinned down uniquely). It is, of course, worth mentioning that there are cases like these, but I want to point out that the most compelling cases for explicature (those that allow one to argue for pragmatic intrusion into truth-conditional meaning) are not of this type. In connection with those cases, I argued in Capone (2006) that they cannot be cancelled.

applies to explicatures as much as to implicatures. And this consideration of regarding explicatures as uncancellable is IN ADDITION to the other considerations which we advance for saying that explicatures specifically cannot be cancelled. (From Burton-Roberts, personal communication). (See developments and limitations of this argument in Capone 2009a)

2.2 The connection with pragmatic intrusion.

Suppose it is accepted that explicatures are cancellable. Then we have problems in explaining how it comes about that explicatures are part of the truth-conditional content of a sentence. Surely, one of the advantages of claiming that an inference is truth-conditional is that it is an entailment-like element of meaning (and, thus, non-cancellable). If we do not want the content obtained through pragmatic intrusion to be entailment-like, why should we make so much fuss about pragmatic intrusion at all? (see also Jaszczolt 2009a in support of the idea that explicatures are entailment-like). After all, we could have a very orderly picture like that of Cappelen & Lepore (2005) in which what is said is mainly associated with linguistic semantics and the various contextual phenomena which we claim to be part of the truth-conditional content of a sentence could be easily assigned at the level of the utterance. Is this what we want? Presumably a radical contextual claim is one that assigns constituents derived through pragmatics inside the sentential content (of course pragmaticians may differ on specific details of their theories; for example, in Capone (2006) I proposed cases of pragmatic intrusion alternative to the standard ones, which in my view were easily assailable). A story like mine, which considers explicatures non-cancellable is more coherent with the view that pragmatics provides constituents of thought intruding into the sentential level of meaning, since entailments are not cancellable unless one cannot deny them. One of the classical tests for entailments is the one by Strawson: if S entails P, one cannot utter S and say not P. So now my question is, do we have enough courage to argue in favour of full pragmatic intrusion? If we have, then we should accept that explicatures cannot be cancelled.³

Of course, the question is rather complex. My opponents may want to insist that explicatures are cancellable by providing examples like the ‘He took off his boots and got into bed’. In Relevance Theory, the explicature is ‘He took off his boots first

³ Responding to an important comment by a commentator, I found it important to say that explicatures are entailment-like, even if not exactly entailments.

and then got into bed'. But this explicature can be cancelled, as in: 'He took his boots and got into bed, but not necessarily in that order'. The same can be said of 'John and Susan are married'. The explicature is that John and Susan are married to each other, but it can also be defeated, as in 'John and Susan are married, but not to each other'.

While these objections are reasonable, all they show is that utterance types are associated with certain potential implicatures/explicatures, and these potential implicatures/explicatures are compatible with a number of contexts, including those in which they need to be cancelled (or aborted). However, my claim is that in a particular context in which certain intentions are fixed, one cannot retract those intentions. Consider, for example, a context in which it is understood that my mother took off her boots and got into bed. Both the speaker and the hearer know what kind of person my mother is, that she would never sin against cleanliness, being the typical Italian mother obsessed with cleanliness. Then there is no question of being able to retract the explicature that my mother first took her boots off and then got into bed. The intentions are fixed in that context. In connection with the following example, there are ways to contextualise it in such a way that intentions cannot be retracted. Suppose you are in court. You are being asked questions by a judge who wants to know whether you knew that John and Susan were married. Suppose that what is at issue is specifically whether John and Susan are married to each other. You answer 'Yes of course I knew that they were married'. And in this context, there is no question of being able to retract the inference/explicature, because you know what is at issue and your answer is an answer to that particular question, and not to another question. Of course, I am not denying that a certain utterance type is associated (by default) with certain potential implicatures/explicatures. And these are compatible with contexts in which they can be retracted. However, if a certain explicature is fixed in context, it cannot be undone. (See Capone 2009).

Before closing this section I want to address a point made by Jaszczolt (2009a) on her paper on cancellability and the primary/secondary meaning distinction. Her idea is that a primary meaning (the speaker's strongly intended main message) is not easily cancellable, and, in fact, it is well entrenched. This idea fits in very well with my idea that strong intentionality tends to block cancellability, albeit Jaszczolt allows cancellability in the sense of 'repair' (or correction). Repair, however, seems to me to conflict with the idea of cancellability. A repair is an evident intention to

about a certain unintended message. Secondary meanings (in short, perlocutionary effects) also are strongly entrenched and difficult to cancel. Jaszczolt notes that explicatures can be cancelled, but explicatures that are part of primary meanings cannot. I suppose I need to modify my ideas in order to allow for a notion of primary meanings to be operative in the context of this chapter. As Jaszczolt says, the enrichment or modulation of a logical form may or may not correspond to a primary meaning. I am essentially interested in this chapter in those enrichments which correspond to **primary meanings**.

3. On modularity of mind

In this chapter I adopt a modular picture of the mind. I shall confine myself to only sketching this picture with broad brush strokes. Work on the modularity of mind started with Fodor (1983) and was extended in Fodor (2000). Fodor mainly distinguishes between input systems (e.g. perceptual processes) and a central system. While modular processes are encapsulated, domain-specific, specialized, shallow, fast and obligatory, central processes for Fodor are not encapsulated but draw inputs from a variety of domains.

This strict picture of modularity has been replaced by the massive modularity picture, mainly advocated by Carruthers (2006), Sperber (2005), Wilson (2005), Sperber & Wilson (2002) and Carston (1996, 2002), among others, according to whom the mind consists of myriad modules, each correlating with a certain (specialized) function, whose input domain is highly restricted (e.g. the visual system is for processing visual percepts only, it does not deal with inputs from other transducers). In this picture modules can share resources, especially if they are situated in close neural areas and if they do not perform concurrent functions. The main advantage of the modularity picture is that it explains how the mind can react in such a fast manner to simultaneous inputs coming from different types of transducers and how it facilitates learning in cases in which two or more different types of inputs have to be analyzed simultaneously. Another advantage of modularity is that each module, in so far as it is to some extent insulated from the remaining architecture, can be damaged, without affecting the remaining modules. So, breakdown in a module creates limited damage, since other modules are available to process information coming from the outside. People can have their language system (language module) damaged while keeping much of the remainder of cognition

intact (aphasia); people can lack the ability to reason about mental states (theory of mind module) while still being capable of much else (autism); people can lose their ability to recognize just human faces (face recognition module); and so forth and so on. An obvious advantage of massive modularity is that processing of information does not create an informational bottleneck, given that various tasks can be carried out in parallel. Another advantage is that the modular system allows evolution to add new modules or to tinker with existing ones, without affecting the remainder of the system. A modular mind is therefore capable of evolving and adding new specialized mechanisms in response to environmental challenges.

Unlike most other authors on modularity, Karmiloff-Smith (2010) puts forwards the view that the mind is gradually modularized during development, as a result of the fact that different areas of the brain are more suitable for dedicated mechanisms (more relevant to certain types of processing). So, on this view, starting out with tiny differences across brain regions in terms of the patterns of connectivity, synaptic density, neuronal type, etc., some areas of the brain are somewhat more suited than others to processing of certain types of input. These ideas are interesting, as we may be open to the fact that modules emerge, not as a result of genetic endowment, but as a result of interaction with the environment and with repeatable patterns of experience.

3.1 Modularity and the mind-reading module

Research on pragmatics and modularity of mind has generally focused on the mind-reading module, which has been held responsible for pragmatic inferences. The mechanisms involved in the operation of the principle of Relevance are usually associated with the mind-reading module. I need to remind readers of the kind of evidence generally bearing on the question whether there is (or not) a mind-reading module and whether certain mind-reading tasks are executed within it. The most widely cited study showing that there is a mind-reading module is Baron-Cohen et al.'s (1985) study of **autistic children**. Autistic children are impaired in mind-reading tasks, and they seem to be blind to the notion that other people have minds. Thus, an autistic child will drag his father as if he were a toy. Autistic children, however, seem to be impaired only in the kind of tasks that involve attributing intentions to others, but they need not be impaired in more generic cognitive tasks.

That mind-reading activities may involve automatic, fast heuristics is proven by patients afflicted with **Williams' syndrome**. This disorder results in an average IQ of around 50, combined with linguistic abilities and social skills. People with Williams' syndrome have good abilities for mind-reading and communication but poor general reasoning abilities. This dissociation seems to support the existence, within the mind-reading module, of a sub-module (or more sub-modules) dedicated to fast and automatic inferences concerning a speaker's intended meaning. (See, however, Perkins 2007, for a deeper discussion).

The existence of this sub-module dedicated to mind-reading is supported by another dissociation. Patients with **Asperger's syndrome** have good general reasoning abilities but serious impediments in mind-reading. These people can use general reasoning to compensate for the lacking special-purpose skills. (See Wilson 2005 for a deeper treatment of this point).

For lack of space I cannot expand on these ideas, but I finally refer the reader to Happè, Loth's (2002) important paper on the dissociation between the mind-reading module in connection with actions and the mind-reading module in connection with communication. This is evidence that the Relevance theory module is a sub-module of the more general mind-reading module.

3.2 The Mind-reading Module and fast-and-frugal heuristics

In this chapter, in consonance with Sperber (2005), Wilson (2005), Sperber & Wilson (2002) and Carston (1996, 2002), I shall try to explain mind-reading in connection with pragmatic inferences (inferences arising from verbal behavior) by positing a heuristic principle working in the Mind-reading module that processes inputs in a very fast way and provide shallow inferences, that is to say inferences triggered by procedures that are designed to provide responses to environmental stimuli in real time, without getting bogged down in laborious reasoning. These inferential processes are fast, automatic, and schematic – in other words they follow what Gigerenzer et al. (1999) call fast and frugal heuristics. In connection with this approach, Carruthers (2006) writes:

Most cognitive scientists now think that the processing rules deployed in the human mind have been designed to be good enough, not to be optimal. Given the speed of processing is always one constraint for organisms that may need to

think and act swiftly in order to survive, evolution will have led to compromises on the question of reliability. Indeed, it will favor a satisfying strategy, rather than an optimal one. (Carruthers 2006, 54).

I produce some simple examples of the type of fast-and-frugal heuristics proposed by Gigerenzer et al. Suppose you are required to answer the question: which of two cities is the larger. A simple heuristic can be useful – you can choose the only one of two options that you recognize.

Another heuristic to be used in response to the same question can be the following: you may look first at beliefs about which properties of cities have correlated best with size in the past. If having a top-division team correlated in the past with greater size, then you select the town which has a top-division team, while neglecting the town which does not have one. If none of the two towns has a top-division team, you move on to the next best predictor of size. It has been shown that fast and frugal heuristics are often successful and are not much below the level of success of more complicated Bayesian algorithms.

The nice thing about these fast-and-frugal heuristics is that they come with stopping rules; you know when you can stop and you do not indefinitely process information in order to arrive at an optimal answer. What is good enough will suffice, and you stop there.

3.3 Kasher (1991) on pragmatics and modularity of mind.

Kasher (1991) is to be considered a seminal paper on pragmatics and modularity of mind. In fact, much of the discussions in this paper has been driven by those initial considerations. Kasher distinguishes two types of pragmatics:

Modular pragmatics; Non-modular pragmatics.

Modular Pragmatics is essentially part of the Language Faculty, and is responsible for the understanding and production of basic speech acts, such as assertions, commands, questions. Non-basic speech acts are derivable from these essential speech acts through the integration of information dealt with by central processes.

Non-modular pragmatics deals with conversational implicatures, institutional speech acts, and indirect speech acts. Central processes are responsible for non-modular pragmatics. Both in the case of indirect speech acts/institutional speech acts and of conversational implicatures the speaker and the hearer must have access to information stored in the central system, information pertaining to institutional facts but also to general knowledge. In general, Kasher takes pragmatic interpretation to be a consequence of Kasher's Principle of Effective Means: 'Given an end, one is to choose the action which most effectively, and at least cost, attains that end, *ceteris paribus*'. (Kasher 1982, 32). In Kasher (1991), he writes:

(...) a cognitive system which derives conversational implicatures involves the application, to the output of some linguistic system, of some general central principles of rationality in intentional action. Hence, to the extent that it seems plausible to assume that rationality principles belong to a central cognitive system, it would be implausible to assume that some *domain-specific* cognitive system produces conversational implicatures. (Kasher 1991, 387).

Of course, Kasher's considerations are conclusive only if it is demonstrated that a principle of rationality is embodied **only** in central processes; could it not be the case that there is a dedicated specialized module, like the 'Theory of Mind' module, which embodies some principles of rationality (e.g. the principle of Relevance) implementing them at a sub-personal level, that is to say not in the form of a consciously accessible reasoning?

Aware that the discussion has become very abstract, Kasher decides to use a more empirical method to settle it and uses the notion of cancellability. Conversational implicatures are cancellable, hence he deduces that pragmatic processes are NOT mandatory. The discussion of cancellability, as pursued by Carston's (1997) paper, bifurcates into a discussion of mandatory and of encapsulated processes. So Kasher's considerations could be taken to militate against the modularity of pragmatic processes in that cancellability of conversational implicatures shows not only that pragmatic processes are not mandatory but also that they are NOT encapsulated (thanks to Robyn Carston for noting this important point).

But at this point Kasher's paper reveals obvious limitations. At the time when the paper was written, all pragmaticians – with the exception of Sadock (1978) and,

ironically, Grice himself – showed unlimited faith in cancellability as the hall-mark of conversational implicatures (and, in general, of genuine non-truth-conditional pragmatic processes). Kasher seems to be adamant that his considerations on modularity are correct because they rest on this almost-universally accepted aspect of pragmatic analysis.

But if my considerations in Capone (2003), (2006) and (2009a) contain some grains of truth, cancellability or lack of cancellability must lead towards a different picture than the one painted by Kasher, namely a picture which is more coherent with recent writings by Relevance Theorists on modularity of mind and pragmatics. Both Carston and Sperber & Wilson, in fact, in recent writings have pushed a modular view of pragmatics, the principle of Relevance being part of the Theory of Mind module, a module with its proprietary heuristics. The spirit of relevance theory, which bounds the search of information and makes use of modules on the fly, seems to tie in very closely with my considerations in Capone (2003), (2006), (2009a) and in this chapter.

3.4 Relevance Theorists on modularity of mind.

The main contributions by Relevance Theorists on modularity of mind are Carston (1996, 1997), Wilson (2005), Sperber (2005) and Sperber & Wilson (2002). These positions move away from the position originally held by Sperber and Wilson (1986), roughly in harmony with the position of Fodor (1983) according to whom the language faculty was modular, whereas pragmatics was roughly constituted by central intelligence processes.

Now Relevance Theorists endorse a massive modularity picture of the mind where the principle of relevance is operative within the ‘Theory of Mind’ module. This picture is different from the one by Borg (2004), since the latter is persuaded that pragmatic processes are not encapsulated and, thus, are similar to theory-formation, such processes being open to information coming from all possible stores of information.

A module, according to the massive modularity hypothesis, is a dissociable component of the mind, with an associated data-base and proprietary cognitive operations. Since modules are dissociable and are characterized by specific functions, it is clear that the operations of different functionally characterized modules must be different. As Sherry & Schacter (1987) say, there is a functional

incompatibility between the operations of different modules, since the operations involved in dealing with a certain function (in resolving a certain type of problems) must be very different from those involved in resolving a different kind of problem. There may be an incompatibility between the solutions to distinct problems. The processes of the ‘Theory of Mind’ module conform to fast and frugal heuristics (Gigerenzer et al. 1999) and are characterized by **satisficing** strategies aimed at drawing inferences that are good enough (for the purpose at hand) rather than being optimal (in the sense that one reaches a guarantee that they are fool-proof inferences). They are characterized by stopping rules (like fast and frugal heuristics in general) – when a good enough solution to a problem is reached, the heuristic principle of relevance stops the search for inferences. An interpretation is relevant if it achieves enough contextual effects for the least possible cognitive efforts. Relevance is, thus, a negative function of cognitive efforts and a positive function of positive contextual effects. The principle of Relevance is a heuristic principle, in that it provides effects that are good enough, and thus must be distinguished from an algorithm which automatically produces fool-proof inferences. One way in which Relevance Theory can advance a view of pragmatics that does not depend on central processes (with its characteristic slow processes, associated revisability of hypotheses and unlimited access to a vast knowledge database) is to claim that the relevance theory inferences are

- a. Fast;
- b. Circumscribed (the search for information is bound);
- c. Mandatory;
- d. Underpinned by a module with fixed neural architecture;
- e. Characterized by a typical developmental/ maturation pattern.

From the point of view of this discussion, the most important characteristics are the first three. The inferences associated with RT heuristics are fast, they usually take fractions of seconds; they are circumscribed, since the context the RT principle searches is not unlimited but is bound by the very Principle of Relevance; they are mandatory. It is simply impossible to refrain from having access to a pragmatic inferential augmentation due to the Principle of relevance. While a person may fail to ‘see’ an inference which is triggered by **personal-level** reasoning (à la Grice), a

Relevance Theory-dictated inference is mandatory and occurs at a **sub-personal** level.

To see the difference consider my usual example of a conversational implicature supported by a personal level reasoning. A student of mine talks about the fact that his scholarship is coming to an end. He is reluctant to ask for an extension of the scholarship or for a recommendation for a new one. However, I infer that he wants me to write a reference for him or that he wants me to give suggestions on scholarships for which he can apply. Surely, in order to have access to the right implicature I must have access to some general script: students in need want scholarships; if the student mentioned the fact that his scholarship is expiring, he may have mentioned that with a purpose; thus, in the hope to reconstruct that purpose, I attribute him the tacit intention of asking me for a reference or for advice. HOWEVER, suppose I already know that the student has just obtained another scholarship for the next year, I wonder what his goal may be in uttering his speech and thus I access another script: when a scholarship expires, one usually needs a letter certifying that the money is well spent (of course scripts may vary quite a lot depending on socio-cultural information; but at least one of the scholarships I benefitted from when I was in Oxford had this requirement). So this is the case of an inference in which the search for information can be extended 'ad libitum'. Instead, an example of a mandatory inference, determined by the Relevance Theory heuristic is the following:

(15)

A: Why don't you join us for breakfast?

B: I have had breakfast.

In order to find out whether the utterance counts as a way of declining the invitation, A must enrich B's utterance in a mandatory way, by adding the time specification (this morning); obviously, it is trivially true that one has had breakfast at least once in one's life!

Carston (1996) discusses an example drawn from a lecture by Wilson (1996) that supports the idea that Relevance Theory processes are encapsulated. The example is the following:

(16) No head injury is too trivial to ignore.

The sentence literally says:

We won't find a head injury whose degree of trivialness calls for any response other than ignoring it.

However, the interpretation we immediately access is:

(17) No head injury is too trivial to attend to. This case seems to show that encapsulation is a property of Relevance Theory's inferential processes.

3.5 Capone on modularity of mind.

In Capone (2010f), I discussed the relationship between Default Semantics and the Principle of Relevance within the modular architecture of the mind. I proposed to come to a better grasp of the interaction in question by utilizing knowledge of the issue of modularity of mind. In particular, I analyzed the phenomenon Bach (1998) calls 'standardization' and proposed that once inferences become standardized, they are no longer processed through the Principle of Relevance, given that they can be furnished directly by the Default Semantics archive. I considered potential objections to this idea, based on experimental pragmatics and arrived at the conclusion that merger representations, which guarantee compositionality at the level of the utterance, take into account both Default inferences and modulated effects due to context. Following Horn (2005), I assumed that contextual information may seep into the pragmatic interpretation while the default semantics is considered. I assumed the existence of a module called 'Mind-reading module' within which the inferential work related to understanding utterances is carried out. I argued that the Mind-reading module includes the processes described by Jaszczolt's Default Semantics as well as the pragmatic inferential processes for which the principle of Relevance is responsible. I argued that Default Semantics describes inferences that make use of

inferential shortcuts stored in an archive and also utilizes simple heuristic procedures (fast and frugal procedures in the terms of Gigerenzer 1999). Default Semantics includes the study of the inferential processes explained through the Principle of Relevance and its basic heuristics (e.g. What is not said isn't) are specializations of the Relevance heuristic. Specializations arise in response to **recurrent** environmental problems and consist of heuristics which, though based on the principle of relevance, have formed a distinct sub-component of the theory of mind, and can be characterized as associative. A certain stimulus will trigger a certain output based on these specialized heuristics as a result of semantic associations.

PART II

4. Applications

In this part of the chapter, my aim is to apply the theory of modularity to specific problems connected with the semantics/pragmatics debate, in the hope that a system of connected analyses will throw light on specific issues. Needless to say, the considerations in the following sections will be theoretical. They are not specifically substantiated through empirical investigation; but of course, in a number of papers, I provided empirical analytic considerations through which I induced a number of theoretical considerations. It would be nice if such considerations could be strengthened by theoretical considerations on modularity through the simple method of letting the issue of modularity of mind interact with the issue of the semantics/pragmatics debate. It is clear that there are some common phenomena in need for an explanation. For example, the raw data of the retina are further enriched through the part (module) of the brain specialized for vision. The data of the retina are two-dimensional, whereas vision has a three-dimensional output obtained through a system of (automatic) calculations on the basis of the respective position of objects. (Marr 1982) (see also Recanati (2002, 121) on the analogy between perception and primary pragmatic processes). Analogously, the input from the Language Faculty is raw; as Carston (2002) says, the output of the language faculty is very often underdetermined – in other words it will not suffice to provide a

fully truth-evaluable proposition. The contribution of the language faculty is sub-propositional and needs pragmatic enrichment to be able to be utilized. The analogy can be pushed further: in the same way in which an action would not be possible on the basis of a two-dimension visual output (imagine what it is like to throw a stone at a target (e.g. a bird in a tree) in a two-dimensional perception in which depth cannot be calculated), a linguistic action would not be possible if one could only resort to un-enriched logical forms (in a sense strictly speaking one should not count them as utterances but as attempts at communication). Consider the following:

(18)

A: Go to the shop in Trafalgar Square in Oxford and buy some wine.

The hearer (suppose it is the case of a foreign student who has just arrived at Oxford) cannot execute the action unless he is able to discern that the action of buying wine must be carried out subsequent to the action of going to Trafalgar Square. The action can be executed ONLY if the hearer realizes the order of the actions and from such an order deduces the location where the wine must be bought. Both in the linguistic and non-linguistic action, the enrichment of visual/linguistic stimuli is a prerequisite for action; and this involves some modular capacity. In the following I provide a number of considerations that seem to show that the modular approach to semantics/pragmatics is particularly fruitful.

4.1 What can modularity of mind tell us about the cancellability of explicatures/conversational implicatures. Modules on the fly.

Followers of Fodor (1983) are against a modular vision of pragmatics because they claim that pragmatic processes are not encapsulated. They usually take (see Kasher 1991) conversational implicatures to show that pragmatic processes are NOT mandatory (hence they are NOT encapsulated). Since conversational implicatures are cancellable, this means that whatever reasoning may be responsible for an implicature, it can interact with vast slices of knowledge and can find a piece of knowledge that can cancel the implicature. (Hence the implicature is not mandatory). The picture of pragmatics which the followers of Fodor had in mind, when resorting to data such as the cancellability of conversational implicatures, is an outdated one.

In two papers, in fact, I claimed that explicatures are not cancellable. If such a claim is correct, then it clearly shows that at least some pragmatic processes are modular in that the search for information is bounded by the principle of Relevance. The two main reasons why I thought that explicatures are not cancellable are that **in context strong intentions are not retractable** and that the non-cancellability of explicatures seems to tie in with the idea that primary pragmatic processes deliver truth-conditional meaning. (See also Horn 2005). Explicatures must have the characteristics of entailments – it is not possible to claim P and NOT P at the same time. Making explicatures cancellable would presumably violate this principle of non-contradiction of entailments. The considerations on strong intentionality obviously interact with the issue of modularity. Explicatures are (usually) processed by the RT heuristics, which are embedded in the Theory of Mind module. Thus an explicature is embedded in a representation of an intentional stance by default. I assume that intentional stance is attributed by default in the Theory of Mind module and from this it follows that explicatures are embedded in an intentional stance representation (hence correlate with strong intentions).

The principle of non-contradiction also seems to operate in the Theory of Mind module, given that the possibility of contradiction correlates with strong intentional stance representations. Without intentional stance representations, there can be no principle of non-contradiction. I assume that explicatures are non-cancellable entailment-like elements because, if we allowed the level of what is said to consist of cancellable elements of meaning, the Principle of non-contradiction would have to be abandoned. The Principle of non-contradiction must be stated at the level of what is said, given that sub-sentential semantics cannot deliver truth-evaluable entities, entities that, in other words, can be contradicted or not. Consider:

(19) John said that P and that NOT P.

There is clearly a contradiction here. However, **contradiction can be assessed only at the level of what is said**, not at the level of sub-sentential semantics (the crucial cases are examples where P is replaced by a conjunctive sentence, as in the basic cases of pragmatic intrusion by Carston). It follows that UNLESS we abandon the principle of non-contradiction we had better accept that explicatures cannot be cancelled. (See also considerations in Capone 2006, about the Principle of Contradiction and what is said)

A reason why pragmatics moved back from a modularity view of pragmatics is that pragmatics seems to involve unlimited access to context – hence the ever present possibility of revisability (hence cancellability). Pragmatic inferences look very much like scientific theorizing, for some proponents of pragmatics *qua* central processes (à la Fodor 1983), and must have unlimited access to a huge database (furthermore one that is continually updated). This seems a conclusive reason for holding the view that pragmatics cannot consist of modular processes.

However, some linguists, most notably Wilson (1996) (see also Carruthers 2006) have insisted on the idea of modules on the fly – modules that limit the information that can be searched in order to arrive at a certain interpretation. For example, consider the case of anaphora:

(20) The surgeon was ready to operate on John. And he was ready to be operated.

The module on the fly posited by Sperber & Wilson restricts the database where referents can be searched to the entities mentioned in the module [The surgeon was ready to operate on John] – as conversation proceeds further, further modules are built on the fly, which are very local in nature. This is also compatible with what Levinson (2000) and Huang (2000a,b) say about anaphora. Referents for pronominals are preferably provided by searching for the values of names already introduced in a local domain (the previous sentence, usually) rather than from a wide domain like extra-linguistic discourse.

4.2 What can modularity of mind tell us about top-down effects - an enzymatic theory of modularity

In this section, I want to dwell on the relationship between top-down inferential processes and modularity of mind (Recanati 2002, 2004). A top-down inferential process is one whereby global contextual considerations project a certain interpretation onto a certain sentence. One of the most influential top-down effects is the resolution of ambiguities in context. In the case of an ambiguous sentence, it is possible to associate two different logical forms with the same phonological string. In fact we have two sentences hidden behind the same phonological representation and, strictly speaking it might even be wrong to say that we have one sentence with

two different meanings. The force of context, however, is so strong that we do not really bother to consider two different syntactic structures, but we immediately assign the sentence its most natural representation. This is a typical top-down effect. You do not bother to consider the meanings of the sentence and then decide which one you want to choose.

An example of a bottom-up pragmatic inferential process is, without doubt, the case of a sentence which contains variables that must be saturated in context. A typical example is a sentence such as:

(21) Look at that.

You cannot know in advance what the value of ‘that’ is but you must first consider the semantics (or the character of the indexical) and then look for referents that are suited to it (perhaps creating a module on the fly, that is to say restricting the search for information to a context that surrounds the thing pointed at; e.g. you do not look for objects behind the speaker’s back). The contextual influence is not top-down because the understanding of the salient elements of the context of utterance do not guide you towards the right interpretation; some sentential elements (e.g. a deictic) function as intermediaries and guide the interpretation process constraining the number of entities that could be selected as referents and also selecting some from the visual field where one is authorized to search for referents (gestures typically helping to pick out the right objects). Jaszczolt (2005) leaves it open whether merger representations are obtained through top-down or bottom-up effects. She is clearly right to leave options open. However, the choice to leave things open is not without theoretical cost. Suppose we assume that a certain sentence interpretation is affected by both top-down and bottom-up pragmatic processes. Then we must assume that there are features that allow the sentence to keep its structure (and basic context-invariant meaning) unaltered. This is clearly useful especially if bottom-up effects are applicable immediately after top-down effects. Unless there is the possibility that a certain basic structure remains unaltered after the application of top-down effects, it is not clear how to proceed further towards the saturation of certain variables.

Now it clear that once we accept that merger representations à la Jaszczolt can be obtained through both types of effects, we are already committed to a certain type of

cognitive architecture. We discard immediately an architecture in which modules are connected through pipes, and instead we adopt the bulletin-board architecture of Barrett (2005). According to Barrett we have a bulletin board and inputs travel towards a certain module and switch it on in case they satisfy certain templates (in a lock and key manner). Once an input is accepted by a certain module, the module provides an output, but the input returns back to the bulletin-board where it can travel towards a different module. This architecture is compatible with the idea that practical reasoning takes inputs from different domains, e.g. desires and beliefs; but it is also compatible with the kind of inferential processes we have described as delivering merger representations. An input can be accepted, say, by the Principle of relevance in the 'Theory of Mind' module, then it goes back to the bulletin board and is available for another modular process. It is important that the input be essentially the same throughout all processes. In this architecture, all representations can be scrutinized by all devices. When a representation fails to meet the input criteria of a device, it is returned to the pool unaltered. When a representation does meet the input criteria of a device, it is processed and then re-posted to the same bulletin board for further scrutiny by other devices. Barrett's proposed architecture is interesting. What this implies for the interaction of top-down and bottom-up processes is that a certain representation travels towards a certain device (say a RT device), is incremented, and then sent back to the bulletin board where it becomes available for a different (or the same) device for a bottom-up process. In an architecture where representations travel through pipes, it is difficult to see how a representation can travel back through a pipe and then be sent to a different (or the same) device for a different type of process.

The enzymatic treatment of modularity supports the idea by Jaszczolt (2005) but also by relevance theorists (e.g. Carston 2002) that there is an interplay between explicatures and implicatures such that it is not necessary to postulate a rigid order (what takes input from what). However, it is also possible to adopt a view similar to that by Clark Barrett (2005) according to which once a representation is assigned a tag, this tag is also used to decide what takes input from what. Suppose that explicatures are assigned a certain tag, then it is possible that conversational implicatures have received an instruction to take input **only** from explicatures. However, the real problem is whether explicatures can also take input from implicatures. According to Jaszczolt and Relevance Theorists they can. Strictly

speaking, it might be possible to have a system in which conversational implicatures take input from explicatures but explicatures can in turn take input from conversational implicatures. What is important is that implicatures can take input only from explicatures (and not from implicatures). However, for all this to be possible it must be the case that an implicature is calculable even without first calculating the explicature; in such a case, the implicature can have an influence on the explicature. A **simple** case would be the following: Calculate the implicature even without assigning referents to variables (e.g. pronominals, deictic elements); then use information from the implicature to calculate the explicature. A more **complicated case** would be: calculate the implicature even if a portion of the sentential structure is missing before calculating the explicature; then use the implicature in order to calculate the explicature.

Let us contrive an example to illustrate this. John and Mary are talking about Dickens. Mary is saying bad things about Dickens (he has not accepted her paper because she is the student of Higgins and Dickens is in the habit of rejecting Higgins's students' papers) (I have contrived the example, these things surely never occur in academic reality!). Dickens is approaching them and is behind Mary, almost within reach of her words, and there is the risk that he might overhear Mary's unflattering remarks. Then John says: 'A linguistics & Philosophy editor' (to borrow one of Stainton's (1994) favourite examples). At this point – even before actually reconstructing the explicature, it is possible in theory for Mary to recognize that what John is saying has nothing to do with the conversation and to interpret this (unilateral and unnegotiated) topic shift as an advice to change topic immediately. At this point, after having recovered the conversational implicature, Mary is able (and has all necessary time) to reconstruct the explicature. The explicature could either be "There is a Linguistics & Philosophy editor over there" or "There is a Linguistics & Philosophy editor just behind your back). It would not be possible to decide which explicature is intended unless there is a clear demonstrative gesture (which is not the case here, while it was possibly the case in Stainton's original examples) or some other kind of information that can advert to the actual message intended. At this point Mary can make use of the conversational implicature to **calibrate** (give a more specific content to) the explicature and arrive at the message intended by the speaker. Since John made a unilateral and abrupt topic shift, he wants to warn her of an impending danger. Surely there would be no danger if Dickens was over there in a

far corner of the crowded seminar room; therefore, Mary understands that John means that there is a Linguistics & Philosophy editor behind her back.

Now, surely this is a contrived cases and I assume that in normal circumstances implicatures take input from explicatures and not the other way round. Therefore, the generalization must be that implicatures can calibrate explicatures **ONLY** if it is possible to understand the implicature without the explicature. Cases where this is possible must be rather limited. So, following Capone (2006), I believe that the generalization that implicatures take input from explicatures is fixed is quite robust, while exceptions must be allowed. At the same time we must concede that implicatures can influence explicatures **ONLY** if there is a non-circular interpretative process and it is possible to understand the implicature even without fixing the referents of variables or without recovering some unarticulated constituent of the utterance (the information that completes the otherwise incomplete propositional radical (to use terminology by Bach 1994)).

4.3 Explicatures as finite inferential processes – satisficing inferential strategies

In this section I want to connect non-cancellability of explicatures with satisficing strategies à la Gigerenzer et. Al. (1999). I thus discuss the idea that a stopping rule is needed for the calculation of explicatures, and this stopping rule goes hand in hand with the notion that explicatures project a strong notion of intentionality. It is convenient to distinguish real explicatures from potential explicatures. Potential explicatures, like potential implicatures, are not cases of real intentionality assignment but are hypothetical cases of intentionality assignment. When one deals with potential implicatures and potential explicatures one says something like: On the basis of ‘S’ (the syntactic and semantic features of S) and on the basis of the fact that the speaker had a reason to use S (and did not use an alternative to S), the speaker intends that P, unless F, where F is some proposition to be derived from the assumptions manifest in context. The clause **UNLESS F** focuses on the hypothetical nature of potential implicatures and *explicatures. However, as far as **actual** explicatures are concerned there can be no ‘UNLESS F’ clause. The calculation process in actual explicatures is quickly over and nothing else can be done with it. We can no longer think of it and process it further in the light of other assumptions. The process leading to explicatures should be finite (both in the sense that we usually

do not spend a long time in coming to an explicature, i.e. we do not ponder on the case of explicature for months, and in the sense that when we acquire the propositional elements that are relevant to the case, we process the utterance, come up with an explicature and close the interpretation process for good). An explicature case is not like a reasoning case where the evidence is sifted and then the experts come to a decision and the decision process could be reopened at any moment (say a trial in which new evidence is crucial and may determine a completely different outcome). An explicature case is closed when the communicative exchange moves beyond the next utterance. Thus, the evidence that is relevant to the communicative process is the evidence available at t , where t is some temporal variable that is indexed to the time of the utterance whose explicature we seek to elucidate and is upper-bounded by t' , where t' is indexed to the utterance next to it. By the time u' is uttered, the explicature of u is calculated on the basis of the evidence available **at the moment** and the case is not opened further. Now, this is important. If we sought further evidence and if u were contextualized say at different moments, t' , t'' , t''' , etc. the explicature could very well be different (given that it arises from the interplay of the linguistically expressed assumption and assumptions available in the context). So an utterance could, in theory, be associated with distinct explicatures. In order to avoid this inconvenience, we have to keep the interpretation process finite. And this is in line with the point that an explicature captures a **unique** intention. If the utterance were considered at different moments, different interpretative possibilities should arise (given the fact that different evidence might be available at different moments). But this is not possible since we have said that the intention to be assigned must be unique.

Now these considerations are compatible with what Sperber & Wilson (1986) say about utterance interpretation. They also argue that the interpretation process is finite and instantaneous – in other words, when the best interpretation is obtained the process stops. My considerations on the intentionality of explicatures make these considerations cogent and provide further justification for them. The considerations so far fit in very well with Gigerenzer et al.'s (1999) satisficing strategies. A heuristic procedure needs to provide results that are good enough, not the best possible results (optimal results). It is clear that satisficing heuristics imply very fast processing as well as a **stopping rule**. When a result that is good enough is found, one sticks to it and there is no way back. Satisficing strategies fit in well with the

idea that explicatures are not cancellable. Cancellability, in fact, implies revisability, which has an extra cost (Jaszczolt 2005) and also seems to be at odds with the very notion of satisficing strategies.

Conclusion

In this chapter I have put many threads together, partly utilizing my previous research and showing that those results make sense if considered from the perspective of a theory of massive modularity. In particular, I have tested cancellability of explicatures, top-down pragmatic inferential effects and Stanley's assumption that primary pragmatic processes must respect grammar in the light of independent considerations from the theory of massive modularity. The most compelling conclusion of this chapter is that a distinction that makes sense, one that is advocated by relevance theorists, is the distinction between sub-personal pragmatic processes (which are fast, mandatory and of which the speaker/hearer is unaware) and personal level pragmatic processes, which are not mandatory, are not necessarily fast, and of which the hearer is aware, given that he plays an active role in reasoning which is similar to that played by the scientist in scientific reasoning. Of course, while sub-personal processes are constrained and operate on modules on the fly (databases restricted for the purpose by the Principle of relevance), personal level pragmatic processes have a freer access to knowledge databases, encyclopedic information, scientific theories, etc. In future work, I may be able to show that even personal level pragmatic reasoning may be constrained by the principle of relevance; however, we may need to investigate a module for scientific theorizing in which the principle of relevance plays a role, albeit a module that is not necessarily innate but may be the result of modularization à la Karmiloff-Smith (2010). This chapter opens this new research avenue.

Chapter 4

The attributive/referential distinction, pragmatics, modularity of mind and modularization.

Abstract

In this chapter I deal with the attributive/referential distinction. After reviewing the literature on the issue, I adopt Jaszczolt's view based on default semantics. I relate her view to Sperber & Wilson's principle of Relevance. I argue in favor of the modularity hypothesis in connection with pragmatic interpretations. I also discuss the issue of modularization à la Karmiloff-Smith in connection with default inferences and, in particular, referential readings of NPs. I then reply to some considerations by Cummings and use data from referential/attributional uses of NPs to show that the modularity hypothesis is defensible.

Introduction

In this chapter I intend to discuss the issue of pragmatics and modularity of mind through an investigation of the attributive/referential distinction.¹ In particular, I want to reply to Cummings' (2009) recently expressed view that the processes involved in conversational inferences are not modular, in that they have unrestricted access to a knowledge data base and deductive inferences. She thinks that general intelligence is responsible for pragmatic increments, whether conversational

¹ Bezuidenhout (1997) too considers that the attributive/referential distinction has a bearing on the issue of modularity of mind and pragmatics, even if her conclusions are different from mine.

implicatures or conversational explicatures. In replying to Cummings, I reiterate my views expressed in Capone (2010c) and I further produce evidence coming from the investigation of the pragmatics of the attributive/referential distinction. Intuitively, default referential meanings of definite descriptions seem to be ideal candidates for modular inferential processes, because they are instantaneous, they arise by default, and are relatively encapsulated. I argue that such default interpretations may interact with contextual clues and that, in limited ways, the defaults can be overridden. But even in such cases, inferential processes are pretty encapsulated. We presumably need a notion of encapsulation that is particularly suited and calibrated in view of the special inferential processes that constitute pragmatic interpretations. Encapsulation à la Fodor will not do; yet, there are alternatives to Fodor's view of encapsulation which do justice to the idea that pragmatic interpretations are not like scientific theories, capable of being revolutionized an indefinite number of times; instead, they are finite, heavily constrained processes which utilize information which has been previously been made pertinent (or relevant) through cognitive nets (unlike Cummings, I believe that modular processes throw a net on what information can be processed and utilized; I call this sort of modular encapsulation **net-throwing**, following a use by Cummings (2009)).

In this chapter I argue that referential interpretations of NPs (and in general default inferences) is the result of modularization. I expatiate on the interaction between Karmiloff-Smith's (1992) theory of modularization and the theory of definite descriptions and argue that the inferences available through the default semantics archive are nothing but re-descriptions of inferences originally available though the Principle of Relevance.

1. Keith Donnellan (1966): Reference and Definite Descriptions.

Donnellan discusses definite descriptions such as:

(1)Smith's murderer is insane

and points out, correctly, that there can be two uses of definite descriptions: a) the attributive use and b) the referential one. In the attributive use, (1) can be used to say that whoever is Smith's murderer is insane (the definite description denotes an x , such that x is Smith's murderer and for all y , if y is Smith's murderer, then $y = x$ (Russell's uniqueness condition)); in the referential usage, 'Smith's murderer' is used to refer to what the speaker wants to talk about, what he has in mind, a particular referent.

According to Donnellan, the attributive use is 'essential' and, presumably, the referential use is derivative (this is my inference, given Donnellan's use of 'essential' for the attributive use). In both uses, the definite description has a denotative function.

Donnellan takes denotation to be distinct from reference (or denoting from referring). He provides the following example to illustrate the difference:

(2) The republican candidate for president in 1964 will be a conservative.

Uttered before the elections, it is very improbable that the speaker was speaking about Mr Goldwater, or that he was referring to Mr Goldwater; even if it could be said that the definite description in (2) denoted Mr Goldwater (since he happened to be the republican candidate for President in 1964).

The attributive/referential distinction is not only observed in assertions, but also in **questions** and in **orders**. If one were to ask:

(3) Who is the man drinking a martini?

one could, thereby, ask a question about a particular person, who is drinking a martini (Who is that man drinking a martini) or about whoever is drinking a martini (I know someone is drinking a martini: who is he?).

According to Donnellan one can have not only a bifurcation between attributive and referential uses, but one can also have referential uses coupled with attributive ones. The case discussed by Donnellan is the following. Suppose I am talking about Jones, whom I believe to be Smith's murderer and I say:

(4) Smith's murderer is crazy.

By 'Smith's murderer' I intend to refer to Jones; however, I am not using Jones' behavior in the docks to justify my belief that Smith's murderer is crazy; I simply rely on the belief that whoever murdered Smith must be crazy to justify my assertion. In this case, we have a basic referential plus attributive usage.

Donnellan, in his paper, claims that the attributive/referential distinction serves to point out some weaknesses in both Russell's and Strawson's views of definite descriptions. Russell's views must be complemented, according to Donnellan, by the idea that a definite description does not only have a denotative use, but also has a referential one. According to him, Strawson's view is wrong on two accounts:

- a) He believes that if a definite description fails to refer because nothing fits it, then one cannot have made a true or false assertion (the question of its truth or falsity does not arise);
- b) He believes that a definite expression can have attributive or referential uses in different sentential frames, but he does not allow for the possibility that the same sentence can be (pragmatically) ambiguous. Instead, according to Donnellan, the same sentence containing a definite description can have either an attributive or a referential interpretation (it is pragmatically ambiguous).

Donnellan likens definite descriptions to Proper Names in that, like them, they can have referential uses.

Now how does Donnellan refute claim a)? His famous example is the following. Suppose that I say:

(5)The man drinking a martini is intelligent,

even if it turns out that he is drinking water, and NOT a martini, I may have succeeded in identifying the man, I have referred to a particular man, and of that man I have said that he is intelligent. And this claim turns out to be true, NOT false. While Strawson would say that in this case we are not confronted with a false or a true statement, according to Donnellan, the speaker has said something which counts as true.

Donnellan also discusses Linksy (1971) who considers the following example:

Said of a spinster, her husband is kind to her, the speaker may well refer to someone. Still the statement is neither true nor false.

Donnellan agrees with Linksy that in case the definite description does not fit the referent, but nevertheless succeeds in referring to it (hence a presupposition of existence is not satisfied), the speaker cannot refer to someone in particular. However, he finds the claim that the statement, made saying 'Her husband is kind to her', is neither true nor false more controversial. Donnellan claims that in this case what the speaker said is true, however we are reluctant to agree with the statement that 'Her husband is kind to her' because there is a **convention of use** prescribing that if someone uses a definite description to speak about a referent, then he should use one that fits the referent. So, the problem we have in saying that the statement that 'Her husband is kind to her' is true is not a question of content, but a question of form. However, we may be ready to say that the speaker said something true about the intended referent.

The last thing to mention about Donnellan's discussion is that the uses of definite descriptions have possibly two presuppositions: a general presupposition that someone or the other is C (where C is the denotation of the definite description); a presupposition that someone in particular is C (only the case of the referential use). These presuppositions are added to the common ground, even if the speaker does not really believe them (Referential use without the belief that the referent fits the description).

1.1 Searle on the attributive/referential distinction

Searle (1979) discusses Donnellan's attributive/referential distinction at length in the light of his own views about speech acts and about referring as a speech act. Searle strongly assimilates the cases discussed by Donnellan to his considerations on speech acts, especially on indirect speech acts. In the same way as, for an indirect speech act, he distinguishes between a primary and a secondary speech act (or illocutionary force) here, in the case of the attributive/referential distinction, he distinguishes between a primary and a secondary reference. The primary reference is the act of referring to an entity through some aspect which may not coincide with the aspect explicitly expressed by the definite description. Thus if, by saying 'Smith's murderer' I am referring to Jones, my statement 'Smith's murderer is insane' is true just in case the predication 'insane' correctly applies to the reference. If a definite description expresses some aspect which is different from the one actually used in referring to an entity, then the aspect encoded by the definite description is secondary.

Searle does not admit that there is a genuine pragmatic ambiguity (attributive/referential). What happens in the attributive use, according to him, is that the aspect under which the definite description refers is primary, rather than secondary, whereas in the case of referential uses, the aspect under which the reference is established is the primary one, not the secondary one (and the linguistic aspect used in the definite description is a secondary aspect). Furthermore, there is, according to Searle, no interesting attributive/referential pragmatic ambiguity because the attributive uses, according to him, are also used to refer. Searle discusses the case used by Donnellan at length to show that attributive uses denote, but do not

refer, and that denotation must be distinguished from reference. The case in question is the following:

A speaker who has uttered:

(6)The republican candidate for President in 1964 will be a conservative

did not intend to refer to Mr Goldwater, even if it happened to be the case that the republican candidate for President in 1964 was Mr Goldwater.

Searle deals with this cogent example saying that, in fact, in a sense, it would be true to say that the speaker referred to Mr Goldwater, even if, due to principles of pragmatics, one does not freely substitute an NP with another in an intensional context. The only reason why we cannot make this substitution has to do with communicative principles, rather than with semantics. Searle shows that in some cases one does feel free to make similar substitutions in intensional contexts, but in others one does not due to facts about context and practical interests.

Searle's account of the referential/attributive distinction seems to have several faults. First of all, he semanticises some facts which, in Donnellans' intentions, were intended to be pragmatic. While Donnellan never explicitly says that, in case a definite description does not fit the referent, but a reference is nevertheless successfully established, the statement is true (he simply confines himself to the more modest claim that what the speaker has said is true), Searle explicitly writes about a statement which is true despite the fact that the definite description does not fit the referent (under the secondary aspect), because the definite description satisfies the referent under a primary aspect. One further problem is that he seems to have drawn analogies from his theory of indirect speech acts, and thus seems to overemphasize the cases in which a definite description does not fit the referent. He seems to make this case standard, whereas he seems to relegate the case in which a definite description identifies a referent in virtue of its semantics 'secondary' (he explicitly writes about a secondary aspect being associated with the linguistic expression (the NP)). The most controversial idea, in my opinion, is that attributive uses are used to refer as well and that there is no distinction between denoting and

referring. Searle apparently takes issue with all those who believe that referring means having something in mind, and these are a great many (see Wettstein on having in mind). Furthermore, he extrapolates facts from indirect speech acts to give the impression that, in using a definite description, such as ‘Smith’s murderer’ one means something like ‘Jones’, one has a primary aspect in mind which is, let us say, an inferential augmentation and is part of the statement made. Thus one when one says ‘Smith’s murderer’, this amounts to saying ‘The man over there’ and he makes it appear that ‘The man over there’ is some kind of unarticulated constituent. (Searle talks about a definite expression’s expressing a primary aspect on various occasions, in the paper, and this gives the impression that the primary aspect is part of the statement made, a sort of unarticulated constituent).

All in all, it seems to me that Searle’s discussion of the attributive/referential distinction is not a considerable advancement with respect to Donnellan’s views, although it certainly points in the direction of inferential theories of the attributive/referential distinction. Crucial to all such pragmatic theories is the fact that there is a distinction between attributive and referential uses, a point which Searle wants to refute.

1.2 Wettstein (1981) on the attributive/referential distinction.

Wettstein also believes that the distinction between attributive and referential uses (of definite) descriptions can be supported. However, he objects to Donnellan’s idea that one can support such a distinction with considerations on what happens when the definite description fails to fit the referent. He thinks that while it is clear that in cases of attributive readings, a statement is false (alternatively neither true nor false), it is controversial that in the case of referential uses, the statement (made) is nevertheless true.

Leaving aside the controversial aspects, Wettstein argues that the case of referential uses is supported by considerations about indefinite definite descriptions. He argues that Kripke (1977) is wrong in thinking that the truth-conditions of the referential reading is captured by Russellian semantics, because the Russellian semantics cannot account for what is being communicated in context through a definite description. Since Strawson’s influential critique, it has been known that the

Russellian truth-conditions for definite descriptions are NOT sufficient to account for communicative uses, since the uniqueness condition notoriously fails in most cases of ordinary uses. When we say ‘The book is on the table’ there is nothing in the sentence that can allow us to pick out a unique table. Defenses of Russell along the lines of elliptical completions of the definite expression, according to Wettstein, fail because on each use many completions are available and one does not know how to choose them; neither is it clear that the speaker must have a completion in mind (he may simply have a demonstrative reference in mind). Wettstein, thus, believes that in most uses of definite descriptions the speakers’ intentions in referring to a certain entity are settled by contextual clues (usually a demonstrative gesture). Wettstein notes that Donnellan’s account of referential uses of definite descriptions is very much in line with this contextual perspective, in which reference is established demonstratively or, in any case, given the rich clues of the context. Furthermore, Wettstein goes on to argue that even attributive uses, which can be accounted for, apparently, through the Russellian truth-conditions, show problems similar to referential uses, in that very often definite descriptions are incomplete and one must resort to demonstrative reference in order to fully specify the attributive reading (The murderer -→ Smith’s murder).

Summing up, Wettstein is able to provide a very cogent and reasonable defense of the attributive/referential distinction.

1.3 Nathan Salmon’s reply

Salmon (1982) takes issue with Wettstein’s treatment in that, according to him, Wettstein’s approach amounts to a defense of the (semantic) ambiguity thesis. Salmon proposes to distinguish between the speaker’s meaning and the sentence meaning. He claims that referential uses are nothing but cases of utterer’s meaning and that both the utterer’s meaning and the sentence meaning should converge and have a common logical form. Salmon reminds us of the fact that it is not uncommon to find cases in which the sentential meaning and the utterer’s meaning diverge, even if one predicts that the utterer’s meaning is a development of the sentential meaning. In particular, Salmon claims that in both referential and attributive uses, the

attributive reading, to be expressed semantically along the lines of Russellian truth-conditions for definite descriptions, is the common denominator.

Now, while Salmon's claims are clearly motivated, it seems to me that Wettstein's view cannot be the real aim of his attack, simply because even if Wettstein says that the referential/attributive distinction is of semantic significance, this does not amount 'ipso facto' to embracing a semantic ambiguity view. It is clear that the view which Wettstein defends is an underspecification or underdetermination view, something along the lines of ideas which became fully explicit thanks to the work of Wilson & Sperber (2002), Levinson (2000), Carston (2002), among others. Wettstein's view is simply that, in most cases, definite descriptions do not uniquely identify a referent and, thus, the Russellian uniqueness condition would not be satisfied UNLESS pragmatic intrusion is allowed. In this sense only, the attributive/referential distinction is of semantic significance. What is meant by that (or what should be meant by that) is that no Fregean proposition can be expressed unless pragmatic intrusion is granted at the level of the interpretation of the definite description, which often happens to be referential.

Summing up, it should be clear that, up to some point, there should be agreement between Salmon and Wettstein, and Salmon's position at least serves to clarify Wettstein's position as to what is the meaning of the statement "The attributive/referential distinction is of semantic significance".

There is a point, which is not without theoretical interest, expressed in Salmon's paper. He claims that while Wettstein easily jumps to the conclusion that even attributive readings have (partially) a referential interpretation, when a complex NP includes tacit or implicit material such as a pronoun, which must be interpreted by reference to contextual information or a demonstration, this conclusion can be avoided along the following lines:

The semantic content of a sentence such as 'His murderer is crazy' is some general proposition to the effect that the murderer relevant to a certain situation as delineated in the context is crazy. I see no compelling arguments against this position.

A more controversial example, discussed by Salmon is the following. Consider the utterance:

(7)The murderer is Jones: Jones is the baby-sitter; the murderer and the babysitter are one and the same person.

Salmon considers that if Wettstein's ideas about the referential uses are accepted, then on the referential interpretation the attributive reading would not be relevant or applicable. But then (7) would have to mean that Jones is Jones. There are problems also on the attributive reading, since on this reading (7) would mean that whoever is the murderer is whoever is the babysitter. But this cannot be the intended meaning. The problem which Salmon does not recognize is that (7) is actually a good case for defending Wettstein's views. Wettstein's view is essentially that a definite description often requires completion (the completion being provided through contextual clues). What is the intended meaning of 'The murderer and the babysitter are one and the same person'? Obviously it cannot be 'Jones is Jones' and it cannot mean 'Whoever is the murderer is the babysitter'. However it can mean 'Whoever happens to be the murderer at t happens to be babysitter at t'. This is a contingent truth; NOT a necessary truth. This is more than enough to support Wettstein's under-determination hypothesis.

1.4 **Kent Bach and the attributive/referential distinction.**

Bach (1981) paves the way for a pragmatic treatment of referential interpretations of definite descriptions. He assumes that referential interpretations too, like attributive ones, have a Russellian semantics, but adds that in addition to this basic semantic interpretation, one further layer of interpretation accrues because of the contextual determinations of the speech act. He accepts a position similar to Wettstein and claims that referential uses are akin to demonstrative reference. In uttering a referential use of 'The F', a speaker, according to Bach, will think of the referent

under some aspect or mode of presentation which may or may not coincide with ‘The F’ and expect the reader to think of the referent *d*, which the speaker intends to refer to, regardless of the mode of presentation through which the hearer thinks of the referent. Bach’s considerations are a consequence of his distinguishing between the sentential meaning and the speaker’s meaning. Bach considers that one should not be surprised if the speaker’s meaning substantially (and systematically) diverges from the sentence meaning. His considerations about referential uses of definite descriptions simply mirror his considerations about the pragmatics of indirect speech acts. His considerations about indefinite definite descriptions are similar to Wettstein’s, as he also thinks that Russell’s uniqueness condition can be satisfied only if incomplete definite descriptions are somehow completed through contextual clues. Bach does not provide a detailed analysis of the detour from sentential meaning to the speaker’s meaning in terms of Gricean pragmatics. I suspect that his way of thinking of this detour is to expect that contextual clues will direct the hearer towards the right interpretation of the definite description as referring to the referent the speaker intends to refer to. Interestingly, Bach does not think that having a referent in mind can explain referential uses.

Bach is not the only one to accept a pragmatic view of definite descriptions. Other eminent authors are Neale (1990) and Soames (1994). Recanati (1989a) unlike them proceeds in the direction of inferential increments called ‘primary processes’ that contribute to propositional forms. I cannot discuss their views due to space limitations, even if I return to Recanati’s ideas in the last section.

2. The semantic turn: Devitt and his critic.

Devitt (2007) develops an anti-inferential or conventionalist account of referential readings of definite descriptions. Unlike Neale (1990) and Bach (2004b), he does not accept that the transition from a quantificational reading to a referential reading is necessary. Instead, he proposes that there is a convention (of use), whereby, by the use of a definite description, the speaker intends to establish a causal/perceptual link to an object. Devitt opposes the particularized implicature view (of referential uses definite descriptions) on the grounds that, according to him, it has no psychological plausibility, given that the preferred standard reading of definite descriptions is the

referential one. He also opposes a standard, generalized implicature view of definite descriptions because he thinks it is simpler to posit a convention for the referential interpretation of definite descriptions. The main reason why he opposes the standard implicature view is that, according to him, in this case the implicature, if there is one, has become frozen, conventionalized. A further reason for opposing the generalized implicature view is that it presumably rests on the assumption that the quantificational reading of the definite description is a route towards the referential interpretation (it should play a role in the calculation of the implicature, even if this role has no psychological plausibility). In addition to this, Devitt argues convincingly that the uniqueness condition can also obtain through his convention for referential usage. Furthermore, Devitt argues that a crucial problem for the quantificational reading is that the uniqueness implication can be applied only through recourse to contextual clues (see Wettstein). A further problem for inferential views like Bach's or Neale's is that they do not make explicit the inferential transition from the quantificational to the referential reading.

In short, Devitt opts for a conventional interpretation of referential uses, which paves the way for a semantic ambiguity, which must be resolved somehow. And how? It appears that pragmatics must be involved somehow in resolving this pragmatic ambiguity. A view that promotes the referential reading by default is preferable. Devitt's view makes some way towards a theory of default interpretation of definite descriptions but does NOT go all the way up this road.

Bontly (2005) replies to the various points made by Devitt. His counterarguments are quite interesting and can be shared to a large extent. From this discussion, I extrapolate a point which can be used to advance the main ideas of this chapter further. While Devitt argues that referential interpretations of definite descriptions are grasped (by the hearers) intuitively, without the hearer having access to a complex reasoning taking into account the Cooperative Principle and the fact that q is required to show how the sentence uttered P follows the cooperative principle. The fact that the hearer has direct access to the referential interpretation, for Devitt, is quite suspect and seems to militate in favor of the idea that there is a convention of

use, a regularity of use, whereby by using a definite description one is actually making reference to an entity (available in context). The reply by Bontly to this argument is that standard implicatures are, in fact, cases in which one does not have access to a conscious argument, to a complex reasoning invoking the Cooperative Principle and the fact that q is needed to make P adhere to the expectations about the Cooperative Principle. These are, in fact, cases in which one has direct access to the inference because the inference has become routinized, it has become standardized. Habit takes over and the implicature becomes intuitively grasped, the inferential process is short-circuited by weight of precedent. Bontly says “Crucially ...the default interpretation remains a conversational implicature; the interpretative habit stems from one’s having calculated such implicatures in the past ...” (p. 8). This is an idea which turns out to be quite useful and which is in line with my considerations in Capone (2010e).

Now, while surely, as a critique of Devitt, the paper is quite persuasive, overall, its most evident weakness is the lack of a clear and reasonable explanation of how the conversational implicature is calculated. Bontly’s main explanation seems to be that in a context in which it is evident that both the speaker and the hearer know that the denotation singles out a particular referent, then that referent is what is intended to be talked about, that referent is what receives a predicate attribution. While surely there is the possibility that the speaker and the hearer, who are capable of identifying the referent through a descriptive condition (the one expressed in the definite description) do so in the course of conversation, it is not obvious to me that this is the right kind of explanation for a standard, generalized implicature. We would expect that a generalized implicature is in general calculable independently of the rich contextual clues. But in this case rich contextual clues are relied on, with the difference that one draws generalizations about what happens or should happen in a context in which the speaker and the hearer are able to identify the referent through the descriptive condition of the definite description. If I am right, the generalization seems to be that in all contexts in which the speaker and the hearer can proceed from an attributive use to an identification of the referent of the definite description, then the description is interpreted as being associated with a referential use. But for me the explanation must be the other way round. If one knows that this use of the definite description is referential, then one will go on to identify the referent. In fact, even in a context in which the speaker and the hearer can easily identify the referent,

it is possible or probable that they will not opt for a referential reading, given that the attributive reading is what they have in mind or is more salient. Suppose I see the corpse and my hearer sees it too, and we both think we know that Jones is the murderer, because he is our only strange neighbor, nevertheless I could exclaim in a raged tone: ‘Smith’s murderer is crazy; look at how the body was mutilated’ and my tone of voice will advertise my intention of talking about whoever is capable of mutilating the body of poor Smith. (See also Donnellan’s case of a mixed referential/attributive usage). Surely here the choice is hard between a referential and an attributive use, but there are inferential steps which may make the attributive reading more likely (such as consideration of the modality of the crime, which plays the argumentative role of support in connection with the attributive reading. That could not play the same argumentative role in connection with the referential reading, because we know Jones to be crazy anyway, we do not need to support this belief about him, since it is the strongest possible belief (by hypothesis).

3. Relevance theory approaches to the attributive/referential distinction

The first author to address the issue of the attributive/referential description within the framework of Relevance Theory is Rouchota (1992). For Rouchota, attributive and referential interpretations form part of the explicature developed on the basis of partial and fragmentary linguistic input, semantic meaning being largely underdetermined. Thus we have a radical departure from previous pragmatic approaches, which were closely related to Grice’s views about conversational implicatures. It is true that Grice considered questions of reference and of ambiguity resolution as part of the proposition expressed, but he relegated other important phenomena which are of significance for propositional content to the status of conversational implicatures, things implicated above the content expressed (or said). Referential and attributive uses of definite descriptions are clearly part of the propositional content of the utterance. However, while Rouchota recognizes that the referential and attributive readings are part of an explicature, it is not clear that she provides a pragmatic derivation similar to that of conversational implicatures. Instead, she treats definite descriptions as if they were similar to pronominals in the referential usage and, like pronominals, capable of being saturated by information derivable through contextual clues. The attributive reading is similarly obtained

through rich contextual clues. In short, there are contexts promoting the referential reading or the attributive reading. This is a heavily contextualist view, which does not take into account the possibility of generalized implicatures. While I and several other authors (e.g. Jaszczolt and Devitt) believe that the referential interpretation is standardly preferred, Relevance Theorists make nothing of this strong intuition.

An advantage of this approach is to show that the length of the definite description may be connected with further implicatures. For example, a speaker who says “The notoriously moody tennis player gave signs of his bad temper when he threw his racquet at his opponent’s head” may implicate that he disapproves of the intended referent, McEnroe, or a speaker who says ‘The fat customer is sitting in his usual chair’ may well convey sarcasm. Consider now Rouchota’s example ‘The man drinking the martini looks miserable’. She considers that the choice of the considerably longer definite expression instead of, say, a demonstrative, must have large cognitive effects, in order to justify the cognitive costs incurred. Thus an implicature may get through that the man is drinking a martini because he is miserable. I think this is a contorted explanation, although it has a grain of truth. If a definite description was preferred to a demonstrative, there must be a reason. This might have to do with politeness, given the precept that one should not point to people especially if they can notice that one is pointing to them. It is simply impolite to point to people, because it is an obvious way of showing that one is talking about them and that one does not care whether other people notice that one is calling attention on them. Another interesting case discussed by Rouchota is ‘Napoleon is in bed’ where one uses a proper name which does not apply to the referent meaning something like ‘The man who believes he is Napoleon is in bed’. This is an inverted commas interpretation. There are interesting remarks in this paper, one of these being that there must be heavy contextual clues to justify an attributive reading. A man who shows surprise at the way Smith was murdered and mutilated, may well say ‘Smith’s murderer is crazy’ without having someone in mind, meaning that the predicate applies to whoever is the murderer. Since he does not know who committed the crime, he cannot have someone in mind. Even if he had someone in mind, yet heavy contextual clues would militate in favor of an attributive reading. Suppose, in fact, that a further contextual effect is to strengthen the proposition that all murderers are crazy. Then such a strengthening would justify the attributive reading.

As I said, despite the interesting things Rouchota says, she does not address the important issue of default interpretations.

Bezuidenhout (1997) makes the interpretation of definite descriptions heavily context-dependent. She explicitly says that while the level of logical form is obtained through (interpretative) processes which are modular in nature (operations pertaining to what Chomsky and his followers called the ‘language’ module, the pragmatic interpretation of definite descriptions is obtained through **non-modular** processes which have access to encyclopedic knowledge and to various types of information coming from the context. She almost makes it appear that the interpretation of definite descriptions is like the saturation process involved in the pragmatic interpretation of pronouns. While surely there may be differences, Bezuidenhout stresses the analogies. Bezuidenhout opts for the underdetermination view of the meaning of definite descriptions and claims that pragmatic information will determine a referential reading in one context and an attributive reading in another. Her semantic view is based on Kempson’s view that definite descriptions activate procedural meaning and that the definite article signals a procedure whose final phase is the recovery of a referent which is accessible in context. The underspecified semantics which Bezuidenhout provides for ‘Smith’s murderer is crazy’ is the following:

MDD: Feature G is instantiated uniquely/accessibly by an x, which is F.

Bezuidenhout interestingly points out that cases which are apparently cases of referential interpretations can also lead to attributive interpretations, as in

(8) Sign: You are entering the Grand Canyon;

(9) Bill Clinton: The Founding Fathers invested me with the power to appoint supreme Court justices;

- (10) Encountering a huge footprint in the sand: He must be a giant.

According to Bezuidenhout 'You' means 'Whoever is entering this place', by 'me' Bill Clinton intends to say 'whoever is the president' and 'He' means 'whoever made the footprint'. The reader is reminded that similar cases were pointed out in Donnellan's paper.

Bezuidenhout's view is based on Relevance Theory because she too believes that meaning is largely underdetermined and that pragmatics serves to make the proposition intended explicit and she also believes that context plays a pervasive role in interpretation, given that the speaker must always make assumptions coming from background information relevant to the interpretation process.

What is not clear is to what extent Bezuidenhout's view is different from Neale's quantificational analysis. Neale's view of the semantics of definite descriptions is that the quantificational reading is at the basis of the referential reading. But it seems to me that MDD is nothing but a different way of saying that a definite description is assigned a semantic interpretation which is that of a quantifier and which also includes a uniqueness condition and a procedure pertaining to accessibility. The only difference I can see is that Bezuidenhout adds a procedure pertaining to accessibility and that she writes about explicatures, rather than implicatures.

Powell (2001) is another interesting article written in the framework of Relevance Theory. Powell discusses the literature on the attributive/referential distinction and claims that the issue of whether the attributive and referential interpretations constitute different propositions (having different truth conditions) must be disentangled from the issue of whether a definite description like 'The murderer' is semantically ambiguous. Furthermore, Powell neatly distinguishes inferential approaches like the one by Neale (1990) or Kripke (1977), according to which one

must distinguish between what is said and what is conversationally implicated, the latter being different from what is said, and inferential approaches aiming at the notion of explicature, a proposition to which both literal meaning and pragmatic inference contribute. The main difference between Powell and the previous theorists is that he relies on the notion of **procedural meaning**, a notion he derives from Blakemore (2000), in order to account for attributive/referential uses of definite descriptions. According to him, definite descriptions encode procedures for determining either a referent or a descriptive content. Like other relevance theorists, Powell assumes that whether a definite description has an attributive or a referential reading must be settled in context. He says that in a context in which the referent that satisfies the description is known, then the referential interpretation comes for free without extra processing efforts. It appears that processing efforts will be essential to the calculation of referential interpretations, assuming that the referential interpretation has greater contextual effects in such contexts. The derivation of the attributive interpretation seems to require for Powell a calculation on the basis of possible alternatives. Given that a directly referential expressions could be used but was not used, then the referential interpretation is automatically eliminated and the attributive interpretation is the one which has most contextual effects. This argument, however, does not stand, because even in the case of referential interpretations one could consider alternatives in which a directly referential interpretation is used and one, could, therefore eliminate the referential interpretation on the grounds that greater processing efforts are involved. But even the argument pertaining to the referential interpretation does not persuade me, because even in a context in which a referent could be clearly intended, an attributive interpretation could have greater contextual effects in case certain argumentative relationships are established between the utterance and further utterances or implicit assumptions.

4. Jaszczolt on default semantics

While Relevance Theorists opt for a theory which is heavily contextual, in that it takes into account the contexts in which the utterances are made, Jaszczolt (1999, 2005) does justice to the idea that the preferred reading of definite descriptions is the referential one. Her theory of definite descriptions derives from a more general

outlook on the interpretation of NPs, whether in normal contexts or embedded in intentional contexts. As Jaszczolt says:

Although definite descriptions exhibit an ambiguity of use between the referential reading and the attributive one, these two readings are not on a par in processing; the referential reading is more salient than the attributive one. (Jaszczolt 2005, 108).

Jaszczolt discusses the example (11)

(11) The best architect designed this church.

According to Jaszczolt, the referential reading corresponds to the utterance that is accompanied by the mental state with the default, strong, ‘undispersed’ intentionality. In the case in which the hearer mistakenly thinks that Sagrada Familia was designed say by Christopher Wren, the intentionality is dispersed as it reaches the object that was not intended by the speaker; likewise, if the speaker falsely believes that Simon Guggenheim designed the Sagrada Familia, the intentionality is dispersed between the intended person and the object recovered by the hearer.

In general, Jaszczolt thinks that NPs strongly correlate with referential interpretations, as shown by her considerations on belief reports, for which she states that the default interpretation is ‘de re’: the believer is taken to believe a proposition about a certain referent.

PART II

5. Modularity of mind

In this chapter, I will not adopt modularity à la Fodor, but the notion of massive modularity (Carruthers 2006), which is, however, connected with Fodor's modularity. The basic idea of modularity is that the brain has a modular organization and that each component of the brain is a module, which is related to other modules in the sense that it can take input from other modules or can give input to other modules. We should not think of modules as specific regions of the brain, even if a module corresponds to a certain neural structure. Since modules can share parts, especially if they are placed at interfaces, it would be wrong to locate a module in a certain area of the brain, as this would not do justice to interconnectedness. Modules are dissociable – and this is perhaps their most important characteristic. Dissociability means that if a certain module is damaged (completely or in part), then the remaining modules can still work autonomously and it is even possible that some other module will try to replicate the processes which were going on in the damaged module. So we shall accept the idea of dissociability, but at the same time we shall admit that the human brain is also characterized by plasticity and that even if certain cognitive processes are best implemented in a certain module, one could nevertheless try to replicate them in a different module (albeit the degree of specialization will be lost). Consider what happens when, due to a stroke, a person loses her ability to read or write. Some authors have agreed that the reading/writing module is the result of modularization (Carruthers 2006; Karmiloff-Smith 1992), rather than being an innate module, and that repeated practice has served to shape the reading/writing module, which has then been partitioned off from the module for object recognition. In other words, the reading/writing module is more specialized than the perception module, and although it may certainly share parts with it, it has been partitioned off from that module, forming an autonomous module. When the reading/writing module is damaged, then the patient can still make use of other modules and replicate the processes which were operative in the reading/writing module. Nevertheless, the reading/writing competence will never be totally recuperated, because the processes occurring in these modules can never become highly specialized as required. They can be a shadow of the previous know-how, but never perfectly suited to the specific

task. So now we are encountering another reason for positing modularity. Modules proliferate in order to adapt to the world's complexity and to develop processes that are perfectly suited to the cognitive needs of a human being. The reader will not be surprised to hear that we, humans, are endowed with a double vision system. One system is suited to identifying objects (and surely it also works for referring), while the other system is more specialized for the navigational needs (Carruthers 2006, Perconti 2008). We orient ourselves in motion through the other vision system. The two systems are complementary. One is more suited for object detection, color detection, the grasping of particulars, etc. The other is less sensitive to detail, but can provide data more quickly and is thus more suitable for navigation, an activity for which colors and small detail matter little, and where it is more important to avoid objects in a very quick time.

The output of modules must be fast, because functional specialization has as its aim providing data very quickly for the various purposes involved in an activity. The output must also be mandatory. In other words, given an input, a module will obligatorily provide an output. This is especially useful in a world in which we and other animals must defend ourselves from predators. We need fast and, also, obligatory reactions. My parrot, however affectionate, has no control over impulses such as:

Fly when a bigger animal tries to catch you.

Even if I just wanted to caress my parrot, his modular processes and the hierarchies of its operations will determine its flight, despite the fact that he knows well that I love him and he loves me, despite the fact that he devotes most of his time to convey his affection (singing, extending his right wing to salute me, turning in circles on one leg, dancing with a right-left and left-right movement every time he sees me). He would fly away from me nevertheless, when I try to catch him, because his brain has been programmed in this way and there is some modular activity which tells him to fly away when a bigger animal approaches him, whoever he is. It is the 'whoever he is' clause that will prevent him from making exceptions. Flying away from dangers

is more important than recognizing who is who, and this is the reason why my parrot has been programmed in this way.

Another characteristic of modular processes is that they are encapsulated. A lot has been written on encapsulation. Massive modularity theorists have weakened Fodor's encapsulation constraint considerably. Encapsulation does not mean (should NOT mean) that a module has no access to another module, but that it has NO access to the operations of another module. It cannot see what is happening in another module, but it can see the result of modular operations, in the form of input (the input of a module is the output of another module). Modules are interconnected and, thus, take inputs from other modules and send inputs to other modules. It is instructive to think of modular interconnectedness through the metaphor of enzymatic processes. According to Barrett (2005), modules communicate through a common bulletin board, where the output of a module is made available to become the input for another module. Every time an operation is made, something is added to the input, but nevertheless, the original input is labeled as having at least the same characteristics as it had before. This is particularly useful when we deal with the relationship between literal meaning and explicatures. We need the assumption that literal meaning receives inferential augmentations, but is nevertheless available for other parallel inferences (we need at the same time to assign referents to pronominals through the perception module and to assign explicatures and implicatures; in order to have all these parallel processes, we need processes which preserve structure, rather than radically transformed. Every transformation is effected in such a way that structure is preserved).

6. Encapsulation, default meanings and referential interpretations of NPs.

I have previously said that the preferred interpretation of definite descriptions is referential. This is clearly the **default** reading. The notion of default reading deserves investigation in terms of the theory of modularity of mind. A default interpretation, in fact, seems to have many of the characteristics of modular processes: it is fast; it is mandatory (unless there are heavy contextual clues militating against this interpretation, one cannot but have access to it (for example, the attributive interpretation of definite descriptions is unlikely to be selected UNLESS there are

heavy contextual clues that favor it)); and it is encapsulated, in a sense which I will make more precise. Encapsulation in pragmatics will be of two different types: 1) activation of the inference in a pre-contextual phase; 2) net-throwing on the contextual information available. I will call these two forms of encapsulation Encapsulation* and Encapsulation**. Encapsulation* basically means that you will go for the default interpretation unless there are contextual clues that militate against it. Even if there are contextual clues that militate against it, the presumption in favor of default meanings is so strong that one tends to ignore context to start with, one takes it into account when the default interpretation really cannot fit into that context. Encapsulation* is a sort of isolation of the information available, a recognition that pragmatic interpretation must start with something and that default meanings are the basic building bricks of pragmatic interpretation. One has access to default meanings, in isolation from contextual information. Of course, contextual information is there, before our eyes, but one pretends that it is not there and proceeds in an orderly way. This is the kind of isolation which a scholar imposes on oneself when one chooses to read, say, what specifically deals with the attributive/referential distinction, setting aside say books on anaphora, however interesting they are. Even if one makes connections, one needs encapsulation of some sort, and needs to rank the possible connections: thus I will make groups of books which I intend to read in a certain order. First I will read book on the attributive/referential description, then I will read books on reference, then I will read books on anaphora, then I will read books on propositional attitudes. Each of these processes of reading a type of books is an **encapsulation process**. One deliberately ignores information which might be relevant but not as relevant as the information one is now considering.

This account of default readings is quite compatible with what relevance theorists say about experimental pragmatics. When a certain default inference is made, it must be made compatible with the context. In other words, a phase of situating the inference in context and of overriding it in case it does not fit the context certainly occurs and we must take this into account.

The inference from definite descriptions to referential readings, I said, is quite standard. Following Jaszczolt (1999, 2005), I accept that there is a strong

presumption in favor of referential readings. This is a more general phenomenon of NPs. In fact, Jaszczolt notices that NPs embedded in belief reports (inside the scope of the belief operator) tend to have ‘de re’ interpretations, rather than ‘de dicto’ interpretations. Surely there is a connection between ‘de re’ interpretations and referential interpretations, since what is understood de re, must also be understood referentially (although things are not the other way round).

Now we are at a fork. Should we say that there are cognitive principles of a specialized nature applying to NPs, such that determine the referential interpretation of an NP in a default context? Or should we say that such default interpretations, which are real, objective enough, can be explained by adopting a more general Relevance Theory perspective? If there are default principles dealing with the referential interpretations of NPs, then we must simply expose these principles. However, if there are general mechanisms of inference, we must explain in what ways the default inferences (in particular the referential readings) are obtained.

I have said before that the RT approaches to definite descriptions were quite ‘ad hoc’ and did not explain the general case, although they could perhaps explain how inferences are operative in particular contexts. We thus need a general treatment of definite descriptions which will produce a default interpretation which is referential.

Things stand in this way, I assume. The human mind is geared toward maximizing contextual effects, while minimizing processing efforts. Referential interpretations are standardly more informative, because they serve to eliminate a greater number of states of the world. If an interpretation is referential, we know what the speaker is talking about and we understand what the speaker predicates of a subject as applying to a particular person. Levinson (2000) and Huang (2000) explain anaphoric processes in a similar way. Pronominals tend to develop co-referential interpretations, since these interpretations eliminate a greater number of states of the world. They reduce cognitive uncertainty, if we want to use a more pretentious term. Now, if this explanation is accepted, the assumptions vocalized in Jaszczolt’s default semantics could be said to follow from it. They are special cases of a more general

case. But then why should we bother with defaults? The same parsimony principle (Occam's Modified Razor) which Jaszczolt invokes many times could be used to say that we do not any longer need her more specialized principles (in particular her Default De Re Principle:

The Default De Re principle:

The *de re* reading of sentences ascribing beliefs is the Default reading. Other readings constitute degrees of departure from the Default, arranged on the scale of the strength of intentionality of the corresponding mental state).

Now, there are two ways to respond to such a criticism. One could be to say that the principle whereby an NP is assigned a referential meaning by default originally derived from more general principles of cognition, but has now become a shortcut for the interpretation of NPs. One could even claim that we need a Reference module and that Jaszczolt's principles are part of that **innate** module.

This idea is interesting of course, instantiating the general idea that when the mind needs specialized principles to deal more efficiently with certain types of information, it develops a module that can deal with that type of information.

A more modest idea is that of **modularization**. We posit modules or archives that store information derivable from pragmatic processes as routinely implemented and transform it into generalizations. Jaszczolt's Default De Re principle could be such a generalization.

Is there evidence that the mind can work in this way? Certainly There is. I mentioned beforehand the case of modularization in connection with the writing/reading systems. The writing/reading systems may derive originally from the perception module, but then while these processes became specialized, and were dedicated to a special problem (how to write or read), the specialized information that connected with this ability (the know-how) became modularized and a new module developed. This is not to say that the module dedicated to writing and reading is innate. What is innate is the predisposition to develop such a module, the

neural structures which give hospitality to the module, the principles for partitioning an existing module from a module that is being developed thanks to information coming from the environment. The possibility of connections between the original module and the new partitioned module must also be innate.

What I am saying is that, by learning how to write and read from the environment (our teachers, our parents, etc.), we store this specialized information in a module that is specialized to receiving and string this type of information.

Could this work for referential readings too? Is it possible that they become standardized and that, when this happens, a module for reference is generated through modularization, the interaction between innate resources and data coming from the environment?

I favor the idea that there must be a module for reference which is the result of modularization and that Jaszczolt's Default De Re principle and the principle relating to referential interpretations of NPs (definite descriptions in particular) must reside in this module. This module is not innately built, but is the result of interaction with the environment.

7. Karmiloff-Smith (1992) on modularization: the case of default semantics

I assume that the issue of default inferences and of definite descriptions ties in very closely with Karmiloff-Smith's discussion of modularization.

Karmiloff-Smith substantially alters the picture of modularity à la Fodor. While Fodor believes that there are input systems (e.g. vision) which are modular and which provide input to central intelligence, Karmiloff-Smith argues that development is the key to understanding the human mind. She says:

Fodor takes as demonstrated that modules for spoken language and visual perception are innately specified. By contrast, I wish to draw a distinction between the notion of pre-specified modules and that of a process of modularization (which, I speculate, occurs repeatedly as a *product* of development). Here I differ from Fodor's strict nativist conception. I hypothesize that if the human mind ends up with any modular structure, then, even in the case of language, the mind becomes modularized as development proceeds. My position takes into account the plasticity of early brain development ... (1992, 4).

The modularization thesis allows us to speculate that, although there are maturationally constrained attention biases and domain-specific predispositions that channel the infant's early development, this endowment interacts richly with, and is in return affected by, the environmental input. (1992, 5).

Karmiloff-Smith's theory is a bridge between theories of innatism and theories like the one by Piaget, who argues that the human mind of the infant is a 'tabula rasa' and grants only some domain-general processes like assimilation, accommodation and equilibration. Karmiloff-Smith hopes to salvage aspects of Piaget's epistemology by arguing that there is far more to cognitive development than the unfolding of a genetically specified program (p. 11) Crucial to Karmiloff-Smith's program is the idea of Representation Re-description which involves a "cyclical process by which information already present in the organism's independently functioning special purpose representations, is made progressively available, via re-descriptive processes, to other parts of the cognitive system" (p. 18). Conceptual re-descriptions are what allow the human minds to make connections between domains of experience which, before the re-descriptive process, were unconnected. Karmiloff-Smith provides the example of the re-description of the concept 'zebra' as 'striped animal' which allows the child to make a connection between the animal 'zebra' and the road sign for a zebra crossing. Re-descriptions are of three types: E1 are not available to conscious access and to verbal report; E2 are only available to conscious access; E3 are available both to conscious access and to verbal report.

Karmiloff-Smith applies her ideas to various domains. However, I propose to focus on her chapter entitled ‘The child as a notator’ since these ideas connect closely with what I said before as the reading/writing module, which appears to be a model for my idea of modularization as I intend to apply it to inferential processes of the default type. Karmiloff-Smith’s main idea is that reading/writing and drawing belong to different modules, as they imply different ideas about what should be done. Even small children are able to distinguish between a drawing and an instantiation of writing. Surely the products may not be neatly differentiated, but they would, nevertheless, clearly insist that this is a drawing and that is an instance of writing. So they have clear in their minds what the constraints on writing and drawing must be. They know that writing involves sequentiality, directionality and ‘spatial frequency or periodicity’ (p. 143).

As Karmiloff-smith says:

Indeed, preliterate children differentiate between drawing and writing even if their “drawings” are not much more than circular scribbles and their “writing” wiggly horizontal lines. But they are adamant about the distinction: “That’s a dog” (a circular scribble unrecognizable to anyone but the budding artist) “and that says “Fido”” (equally unrecognizable, but a horizontal squiggly line). (...) Moreover, video tapes show that preliterate toddlers lift the pen much more frequently when pretending to write than when pretending to draw. The toddler goes about in the process of writing and drawing differently, even if the end products sometimes turn out similar. It is essential to distinguish product and process, because the toddler’s notational products may at times appear domain-general to the observer, whereas their intentions and hand movements bear witness to a clear differentiation that they have established between the two systems (p. 143).

Children who have been asked to distinguish between writing and drawing do not confound drawing with notation as they make clear-cut distinctions between the two notational domains. Drawings are rejected for written language and single elements

are accepted for number, but rejected for writing, linkage between elements is accepted for writing but not for number notation.

Karmiloff-Smith rejects the idea that these constraints are innate, while she accepts the idea that there are such constraints and that they are organized in modules (she furthermore says that the modules for writing/reading and drawing are in different hemispheres). Her reason for rejecting innatism and embracing modularization, in the case of the reading/writing module is that reading/writing, unlike verbal production, are relatively recent in terms of evolutionary time. She says that hundreds of thousands of years of evolution were needed for spoken language to become biologically constrained, but the use of cultural tools for writing dates back only 5000 or 6000 years. So it is implausible to invoke an innately specified bias for writing. (p. 147). Karmiloff-Smith argues that they are due to a process of modularization that is the product of learning during childhood.

My necessarily brief treatment of Karmiloff-Smith's ideas on modularization paves the way for my ideas about modularization of inferential processes. In my previous paper on Default Semantics and the architecture of mind (Capone 2010c) I proposed that a default semantics archive is built to store regularities of inferential results. Jaszczolt's Default De Re Principle and the tendency to interpret definite descriptions as referential expressions may be due specifically to the standardization process which short-circuits an inferential process to a cell in a memory system (a default semantics archive) which produces directly the result of the inferential process. We may propose that cells in the default semantics archive do not simply supply the results of inferential processes one by one, but may be organized in principles of a more general nature, such as the following:

For any NP, go a referential interpretation first.

Such principles are the result of modularization, of learning, even if, obviously, they interact with more general principles of cognition such as Sperber & Wilson's

Principle of Relevance, according to which relevance is a positive function of contextual effects and a negative function of processing efforts. I argue that referential interpretations of NPs are default because they obey the Principle of Relevance. In fact, a referential interpretation eliminates a greater number of states of the world, while an attributive interpretation is compatible with a certain number of referents (The President of USA: Clinton, Obama, Kennedy, ...). In my view, fully identifying information is to be preferred to descriptive information (only) and, thus, the referential interpretation prevails at least in those cases in which the hearer is interested in knowing specifically who did the thing in question, who acted in such and such a way.

Furthermore, modularization, according to Karmiloff-Smith, involves re-description and, in particular, re-descriptions of the type E1, E2, E3. In the case of modularization involved in the creation of a default semantics archive, we certainly have the phase of re-descriptions E2, E3, since the default semantics archive allows access both to consciousness and to verbal report. Re-description is clearly involved in the modularization process of referential interpretations of NPs. Before modularization, in fact, an NP must be considered as a semantic structure allowing interpretative ambiguities. Instead, after re-description, interpretative ambiguities are eliminated. Something has occurred to change linguistic knowledge. The NP has been marked as + referential after re-description.

8. Capone (2010c) on pragmatics and modularity of mind

Capone (2010c) is a discussion of modularity of mind as applied to pragmatics. Capone's belief is that a critical discussion of modularity of mind can improve our understanding of the semantics/pragmatics debate. The main points addressed by Capone are the following: 1) merger representations and enzymatic modular processes; 2) fast and frugal heuristics and satisficing strategies; 3) cancellability and 4) modularity; pragmatics and encapsulation.

Concerning the first point, Capone argues that Jaszczolt's theory of merger representations (Jaszczolt 2005) provides the principles of compositionality for acts of communication. Capone accepts that compositionality is best instantiated at the level of merger representations – representations that combine information coming

from different sources (1) semantics; 2) lexical defaults; 3) socio-cultural defaults; encyclopedic knowledge). Capone argues that, since merger representations combine outputs of both top-down and bottom-up inferential processes, a modular view according to which modular connections are NOT pipes must be accepted. Instead, modular connections must be conceived of as enzymatic processes taking input from a common bulletin board and providing output to this very bulletin board. Enzymatic processes explain how outputs of a process can become input to another type of process. Concerning point 2) (satisficing strategies), Capone considers inferential processes (of the unreflective type) as instances of fast-and-frugal heuristics whose aim is NOT obtaining an optimal result, but only a result that is good enough. As soon as a good enough inferential process is instantiated, the search for pragmatic interpretation stops. Relevance Theorists have drawn attention to these inferential processes. These processes interact with our view of modularity because modular processes are also fast, automatic and finite.

Concerning point 3 (Cancellability), Capone argues that one of the main obstacles to a modular view of inferential processes is removed once it is recognized that explicatures are NOT cancellable (a point discussed at length in Capone 2006, Capone 2009; see also Burton-Roberts' 2005 splendid paper). The convergence between Capone's previous work on lack of cancellability of explicatures and modularity of mind is not a mere coincidence. Lack of cancellability supports the view that inferential processes are modular in nature, given that they are not optional, but they are mandatory (cancellability of explicatures threatening the idea that inferential processes are mandatory).

Concerning point 4 (encapsulation), Capone argues that pragmatic processes are unlike theory-formation (whereby a theory is continuously revised). They have access to limited information, which is encapsulated through the Principle of relevance (see the discussion by Capone of modules on the fly). Capone concentrates on inferences which are automatic and belong to the non-reflective type. Much else must be said about non-automatic, reflective inferences, about which Cummings (200) has much to say. We shall now consider Cummings' views in the next section.

9. Louise Cummings, modularity and pragmatics

Cummings (2009) discusses the same issues discussed by Capone (2010c), but arrives at different conclusions. I assume that the main difference between Cummings and Capone's ideas is based on a different focus. Cummings uses examples of reflective inferences, in which a speaker embarks on a reflective task which uses explicit arguments and reasoning in order to calculate the intended point of an utterance. Capone, instead, focuses on unreflective inferences, which are fast and automatic and in which the calculation of the implicature is not available for conscious access. It is not surprising, therefore, that they should arrive at different conclusions. Yet, intuitively both types of inferential processes are operative in non-logical inference, and thus 'prima facie' both authors say something important. Surely, in discussing implicatures, we must take into account both non-reflective and reflective inferences; yet, it should be clear that unreflective inferences have a privileged status, since they are those which enter into primary pragmatic processes in so far as they contribute to the explicatures, to the explicit contents of utterances and thoughts. Considerations about reflective pragmatic processes surely are important, and yet they cannot be used in isolation to prove Cummings' main point: that is, that pragmatic processes are NOT modular.

Cummings discusses two views of modularity: she approves of the former and she criticizes the latter. Cummings accepts Kasher's (1991) idea that there is a pragmatic module which processes speech acts and determines the illocutionary force of an act of communication on the basis of some presumptions which are usually triggered by the syntactic form of the utterance. There are rough correlations between declarative form and the force of an assertion; between imperative form and the force of an order or command; between an interrogative form and the force of a question. These correlations are standardly used to calculate the force of an utterance, but they can be overridden and, thus, in context, the force of an utterance can be quite different from the presumptions calculated in the pragmatic module. A central system receives input from the pragmatic module and determines in context the particular illocutionary forces of utterances.

This view is contrasted with the view of Theory of Mind (ToM) theorists, for whom the demarcation of the Theory of Mind module is obtained by certain restrictions on the flow information between the psychology faculty and other cognitive domains. Segal's (1996) is the most representative voice in the ToM camp:

In particular there may be a one- or two-way filter to information. In Jerry Fodor's (1983) terminology, intentional modules may be 'informationally encapsulated': some of the information of the subject's mind outside a given module may be unavailable to it...And going the other way, intentional modules may exhibit 'limited accessibility'; some of the information within a module may be unavailable to consciousness...I suggest that if a set of appropriately related psychological states exhibits either informational encapsulation or limited accessibility, then they constitute an intentional module (Segal 1996, 143).

Cummings dismisses this important theoretical step by a complex reasoning:

However, for our present purposes, a modular approach to ToM is only interesting to the extent that it can explain certain features of pragmatic interpretations. For example, a ToM module must be able to capture the relative ease with which a listener is able to recover the implicature of an utterance. We will see that these beliefs are not domain specific or established in advance of interpretation, as they would have to be if they were mediated by the processes of a cognitive module. In fact, many of the beliefs that are integral to utterance interpretations are not even beliefs about our interlocutor's mental state (although, of course, many others are). We will also see that beliefs are revised, rejected and reinforced by a whole range of contingencies in the listener's environment and by other beliefs that are stored in memory or that are the product of more general inferential processes (...). It is difficult to see how these various contingencies and other beliefs can even get access to a ToM module, given its informational encapsulation. In the absence of such access, one cannot begin to imagine how a ToM module can capture the cancellability of implicatures and the defeasibility of presuppositions, for example. (Cummings 2009, 161).

The good point that Cummings is making is that a ToM module must integrate different sources of information, some of these pertaining to encyclopedic knowledge, the beliefs of the hearers, and also the logical inferences made online by

the hearers. It is clear that the final result of pragmatic inference must be able to combine different sources of information. The term merger representations invented by Jaszczolt (2005) does justice to this idea that pragmatic inference combines different sources of information. For pragmatic purposes, the main sources of information which must be integrated or merged with others are the results of default inferences and the result of genuine pragmatic inference based on rationality principles. But given that a global pragmatic process must incorporate a less global pragmatic process, is it impossible that at some stage of the pragmatic process some sort of encapsulation to be defined not along the lines of Fodor (1981) but along the lines of theorists of massive modularity (Carruthers 2006) can be posited? At some point in her most interesting discussion, Cummings says that it is not the case that one throws a net on the information that is accessible to pragmatic processes. She thinks that pragmatic processes are permeable to a whole range of knowledge and deductive inferences. In my discussion I will use the terms **net-throwing** and **permeability** as key terms of the discussion of genuine pragmatic processes. Net-throwing is important, because there must be ways to limit the information that must be considered in calculating a conversational inference. Net-throwing mainly consists in obeying the Principle of Relevance which applies not only to the calculation of inferences, but also to the provision of contextual information to the inferential process. Context does not provide an unrestricted flow of information, but provides a restricted flow of information. The only information that goes through the net is relevant information, information which has a promise to interact with linguistic information in interesting and fruitful ways, by incrementing it and making its meaning potential optimal. Furthermore, contrary to what Cummings and Kasher accept, the processing of pragmatic inference is **UNLIKE** theory-formation. In theory formation, theories are developed, challenged, revised, and the process is possibly infinite. Instead, utterance processing must occur in real time, and is usually over when the next utterance occurs and the conversation flows in the direction of what is said next. There is no time, usually, to go back and revise interpretations, following the procedure of theory-formation. If theorists use the notion of cancellability to prove their point that pragmatic processes are permeable to different sources of information, one would have to reply that, yes, pragmatic processes are permeable to different sources of information, but under the constraint that the pragmatic process is finite, fast, circumscribed to the present. Furthermore, while potential implicatures

and explicatures are usually cancellable, once intentions are fixed in context it will not do to cancel them. This is an idea which I expressed in Capone (2006, 2009, 2010c) (but also see Burton-Roberts's 2005 important work) and which has been supported by Burton-Roberts (personal communication). So, I agree that the pragmatic process is permeable, but it is **instantaneously permeable**, and thus it is completely different from theory formation, which occurs without any time constraints. While theory formation is a collaborative process which involves many scientists, utterance interpretation generally involves two speakers (or a limited set of speakers who are located in space relative to one another). So, the main difference between theory formation and utterance interpretation is that the environment of the conversation provides fundamentally crucial (perceptual) input to the pragmatic processes. Such a constraint is obviously not in place in theory formation, where scientists are usually situated in various parts of the worlds, and the setting does not have a crucial role in anchoring utterances in time and place. The inferential process in ordinary conversation is necessarily circumscribed and encapsulated and the most crucial sort of net-throwing occurring is that due to place and time.

What about the constraint that when optimal relevance is achieved, the interpretation process stops? Is not this a sort of encapsulation? Cummings, in her discussion of schizophrenia, mentions a notion of praeter-relevance as advanced by important work by Cram & Hedley (2005). A problem which schizophrenic patients experience is that they process utterances without stopping when the first interpretation that satisfies the Principle of Relevance is obtained, but they go on making further inferences. If Cummings mentions such a problem, then surely it must be the case that the inferential process is finite and has stopping rules. And this means it is circumscribed and this is surely a limit to permeability. Permeability occurs to some extent, but is not unconstrained. I will thus talk of **circumscribed permeability**, which is still a modular notion.

As I said, Cummings provides an interesting example of reflective inference, in which a number of pieces of information flow into the interpretation process. The example is the following:

(12) Sam: Do you come here often for a walk?

Tom: I hold down two jobs, so what do you think? It's not as nice as it used to be. Owners are letting their dogs foul the pavements and there is litter everywhere. It was local teenagers who vandalized the benches.

Cummings is clearly right in saying that a number of pieces of information flow into the inferential process, as we surely need to take into account certain beliefs such as 'People who have two jobs have little leisure time'. There is no doubt that a number of beliefs must be involved. Yet, it is the principle of Relevance which circumscribes the process and prevents one from resorting to unnecessary beliefs. Furthermore, her example is clearly a case of **reflective inference**, while I said that in order to investigate the issue of modularity it would be best to examine cases of non-reflective, instantaneous inferences.

10. Further considerations on modularity and definite descriptions.

I propose to use the issue of definite descriptions to throw light on the issue of modularity of mind, given that cases of referential uses of definite descriptions are cases of non-reflective, instantaneous, automatic inferences. My view runs counter to what Bezuidenhout (1997) says about the non-modularity of pragmatic processes involved in interpreting definite descriptions. As I said before, default inference is in itself a case of encapsulation, since it requires that for an infinitesimal period of time the inference is encapsulated and processed in isolation from contextual information. Contextual information flows in at a second stage, when the speaker or the hearer needs to assess whether it fits the context or whether it should be replaced with a more tailored inference. Both default inference triggering and contextual tailoring are processes which involve encapsulation. Contextual tailoring involves encapsulation in that the context considered is circumscribed by the Principle of Relevance. Suppose you hear:

(13) The President of the United States arrived in Italy today.

If the sentence was uttered in 2010, you will tend to think of Obama that he arrived in Italy today. There is a default interpretation which presses to become an actual

implicature and the features of the context are such that filter this implicature in. In general, the feature PAST Tense reinforces the default interpretation because it is unlikely that the speaker wants to talk of a past event without having an actor for that event (see Higginbotham 2009). And the actor cannot be whoever is the President, but someone in particular. (13) says that there is a president of the USA, which is uniquely identifiable, and Obama, who is the President, arrived in Italy. We could go on to say that, if there is an actor, then the actor fills the denotation of the definite description, and this will give us the referential interpretation by default. But if this is the case, then the case must be extended to progressive as well:

(14) The President of the USA is flying to Italy today.

Not only does the definite description trigger a referential interpretation by default, but the sentential frame favors this referential interpretation. The sentential frame, by default, can be said to play a role in promoting a default inference. Things are not very different in the future, provided that the time is specified:

(15) The President of the USA will fly to Italy tomorrow.

The hearer is unlikely to interpret the sentence as meaning that whoever is the President of the USA will fly to Italy tomorrow, since ‘tomorrow’ specifies the time at which the event is located, and if an event is located at a certain time, there must be an actor at a certain time.

Consider now an important example discussed by Powell (2001):

(16) The President of the United States changes every five years.

Powell thinks that uses of (16) are meta-representational, because we cannot clearly intend a referential interpretation, nor can we intend that whoever is the President of the United States changes every five years. According to him, (19) expresses a

proposition to the effect that the denotation of the descriptive individual concept corresponding to ‘the President of the United States’ changes every five years. Powell writes:

This interpretation is accessed, as ever, via considerations of relevance: neither the straightforward *de re* or descriptive interpretations achieve optimal relevance, since, for both the construction of a context in which the interpretation yields sufficient contextual effects puts the hearer to too much processing effort. The proposed interpretation, however, yields sufficient contextual effects... (p. 122).

Something along the lines of what Powell proposes must be true. However, it should be clear that the reason why the meta-representational interpretation occurs here is that the sentential frame is different from (15) and (14). We are not confronted with an action located at some point in time, but with a generic sentence. Since a generic sentence has generic validity, which is not limited to a certain period of history, we understand that the sentence cannot talk about a particular president, since it is necessary that Presidents will be different at different periods, and the possible interpretation ‘Whoever is the President changes every five years’ cannot be the right one, because it would presumably say that whoever is President changes clothes every five years’ (pretty implausible, isn’t it?). The meta-representational interpretation goes through because, as Powell says it yields sufficient contextual effects. Given that we arrive at the explicature in order to avoid implausible literal meanings, it is clear that it is not easy at all to defeat the inference, intentionality being fixed by the search for plausibility

Now, in the picture I am proposing, default interpretations are tried first, but if they do not yield sufficient contextual effects, they must be replaced with interpretations that are more tailored to the context. Sentential frames signal, in general, when default interpretations get through or not.

It may be claimed that certain sentential frames promote, instead, attributive readings. Consider, in fact, (17), (18) and (19):

- (17) John wants to become the President;
- (18) John wants to be the President;
- (19) John wants to be appointed professor of Linguistics.

Clearly, these are NOT referential uses, but they are attributive ones. In (17) John wants to have the attribute ‘The President’ (and wants the transition from not being the President to being the President); in (18) John wants to have the attribute ‘the President’ , but no mention is made to wanting the transition from not being the President to being the President. In (19) too John want to have an attributive ‘Professor of Linguistics’.

In none of (17), (18), (19) do we have the interpretation ‘John wants to become Obama’. The referential interpretation which, as we have seen, usually arises as a default, in this case will be aborted because of the general belief that it is impossible to try and be another person, however hard one may try. Even if I wanted to, I could never be Obama. Would it be reasonable to say that in these cases the default is cancelled? Is it not preferable to say that sentential frames like ‘become NP’ also have default interpretations, and these are different from the referential interpretations? This issue would not be, clearly, otiose. However, I cannot settle it here. Consider now the following cases, the first of which is discussed in Higginbotham (2009):

- (20) Heimson believes he is Hume;
- (21) Heimson believes he is the President of the USA.

(20) has a ‘de se’ interpretation and Heimson can believe he is Hume only if he does not believe that he is Heimson (Heimson being different from Hume). It makes sense to utter (20) if we know that Heimson does not think of himself under the mode of presentation ‘Heimson’ but only under a first-personal mode of presentation

(see also Capone 2010 a)². If we replace the proper name with a definite description, we obtain a sentence like (21) whose preferred interpretation is not the referential reading, but the attributive one. In other words, we are charitable enough to use a minimal departure from rationality, and, thus, even on the assumption that Heimson is crazy and believes extravagant things about himself, we assign him the least extravagant belief. Believing oneself to be the President of the United States involves a smaller departure from rationality than believing oneself to be another person (to be Obama, for example). Thus, in interpreting sentences like (21), by default, we adopt the minimal possible departure from rationality and we prefer one of two interpretations, if it is the least extravagant thing to believe. This **charitable** interpretation will be preferred over the least charitable interpretation. We may assume that more charitable interpretations get through by default. So, in interpreting (21), we first of all try the referential interpretation, but then by default opt for the more charitable interpretation. In calculating the explicature we go for plausibility and this is why we are reluctant to give up the explicature and to cancel it.

Now, we have seen that there are departures from default interpretations. This means that default interpretations interact with sentential frames or other contextual assumptions. Does this imply that we must give up the modularity hypothesis? My reply is: in no way! All that the modularity hypothesis compels us to accept is that pragmatic processes are fast, mandatory, encapsulated. In case (21), it is clear that we do not consider different hypotheses and chose one and then are open to the possibility that the hypothesis is revised. The interpretation process is finite and we utilize for this process only information that is relevant. In particular, we utilize information to the effect that people who believe they are the President of the United States are more normal than people who believe that they are Obama. And why do we utilize this piece of information? We do so because it helps us choose between the referential and the attributive interpretation. Given that our interpretative problem is how to choose the referential or the attributive interpretation, we bring into the process extra information, only on condition that it helps us resolve our original problem. So the basic constraint to follow in bringing in additional information is the

² Even this is the result of pragmatic interpretation, since believing to be one person involves a smaller departure from rationality than believing oneself to be two (different) persons at the same time).

principle of Relevance which induces us to maximize information, to look for interpretations which maximize contextual effects, with minimal processing efforts.

We have seen that definite descriptions usually involve referential interpretations. However, Bezuidenhout (1997) has shown that the question of the attributive/referential distinction arises with pronominals too. Bezuidenhout considers cases like the following:

(22) Sign: You are entering the Grand Canyon;

(23) Bill Clinton: The Founding Fathers invested me with the power to appoint supreme Court justices;

(24) Encountering a huge footprint in the sand: He must be a giant.

Bezuidenhout takes (22) to plausibly mean that the visitors are entering the Grand Canyon; (23) plausibly means that The Founding Fathers invested the American President with the power to appoint supreme Court Justices; (24) plausibly means that whoever left the footprint in the sand must be a giant.

These examples are of interest because they instantiate cases in which a pronominal, despite the fact that it has a referential interpretation by default, is assigned an attributive interpretation due to contextual assumptions. But this means that we must assume that pronominals have referential interpretations by default. In some cases, however, these default interpretations are overridden by contextual information. Since we know that it is not the case that the Founding Fathers invested Clinton with the power to appoint supreme Court Justices, we assume that what is meant (the m-intended point) is that the President, whoever he might be, was given such powers by the Founding Fathers. Cummings leads us to believe that the informational increments due to contextual processing which enter into the interpretation are cancellable. Cancellability, according to Kasher (1991) and to

Cummings (2009), attests to the fact that pragmatic inference is a truly global inferential process, like theory-formation. Yet, is it not clear that, when we settle on the reasonable interpretation of (24), the other alternative has to be abandoned and the inference is hard to cancel? So, while inferential augmentations seemed to prove Cummings' point, in fact they prove to be the most thorny cases for the claim that explicatures are cancellable (they also prove to be thorny for Bezuidenhout (1997) who claims that the processes involved in interpreting NPs are not modular). We often resort to explicatures to show that an implausible interpretation is replaced by a plausible one. But it is this need for plausibility which militates against easy cancellation of the inference. And if an inference is not cancellable, then a case can be made for the view that the inferential process which produced it was modular in nature.

Before ending this chapter I would like to call attention to an example discussed by Allan (2010) in an important paper in which he considers reference an act of communication. The gist of Allan's view is that reference is an act of communication that exploits contextual clues utilized by the hearers to establish the intended referent. Being immersed in the theory of Pragmemes broached by Mey (2001) and discussed further by Capone (2005), it is not surprising that Allan should make us see reference as a process that heavily relies on contextual clues. While discussing Jaszczolt's view that "intentionality cannot be called a process, it is an instantaneous 'firing at', 'targeting' objects, it is not an object of passing from sense to the referent (Jaszczolt 1999, 112) and that definite descriptions trigger, by default, referential interpretations, Allan considers a counterexample. Consider:

(25) The best architect designed this church.

According to Jaszczolt this sentence in context means that Antoni Gaudì designed la Sagrada Família because socio-cultural defaults are immediately activated on hearing the sentence. According to Allan, instead, the preferred interpretation in this case is the attributive reading, while he concedes that conversational implicatures are responsible for identifying the referent with Antoni Gaudì. According to Allan, the default interpretation is attributive. The speaker means (something like): 'X designed this church, and he is the best architect'. What should we make of Allan's critical

position? Does it seriously militate against Jaszczolt's view that NPs are referential by default? And does this view militate against my modular view of pragmatics, according to which pragmatic interpretations are fast, automatic, mandatory and encapsulated, at least in the case of non-reflective inferences? An easy answer to this question would be to say that even fast, mandatory, encapsulated inferences are sensitive to contextual information and can evaporate if there are contextual clues that militate against them. However, this is NOT the point of Allan's objection to Jaszczolt. He claims that in this case the default interpretation is attributive. We should go along with Allan if we recognize that various examples are like (25):

- (26) The worst murderer killed Jones;
- (27) The best butcher sold us the meat for the evening dinner;
- (28) The most beautiful actress was chosen for the party;
- (29) The singer who had the best voice was asked to sing at my wedding.

The use of superlative constructions points to a use of the definite description which is not referential but argumentative. By default the utterance is given an argumentative role (support of another statement) and the hearer is driven to a search for the argumentative relation which serves as the glue for the utterances under consideration. Is this a case in which a default overrides another default? It is not impossible to argue that different sentential frames are associated with different defaults. Another possibility must be considered. The more materials we add to a definite description the more likely it is that it will be interpreted attributively. While we predict referential interpretations for definite descriptions on the assumption that the referential interpretation is more informative, adding descriptive materials makes the dimension of cognitive efforts having a greater weight and, in order to counterbalance this weight, the attributive interpretation prevails with the understanding that it has a function in determining the argumentative role of the utterance. It seems to me that the considerations triggered by Allan's ideas on definite descriptions lead us to a position which is very different from the one embraced by Recanati (1989). According to Recanati, in fact, referential interpretations require heavy contextual processing, while attributive interpretations do not (and this is line with his view that definite descriptions are unmarked with

respect to the feature + referential, while directly referring expressions (e.g. proper names, pronominals, demonstratives) are marked with respect to the feature + referential). In my view, following Jaszczolt, Devitt, and also Allan to some extent, definite descriptions associated with certain sentential frames are marked as + referential by default, while other sentential frames are marked as – referential by default, depending on the amount of processing effort involved by the presence of extra linguistic materials. It is also possible to see things differently, the issue being not a matter of having different defaults, but of showing that the addition of further descriptive materials changes the default. The role played by the Principle of Relevance in triggering the search for argumentative relations seems to attest to the fact that the default interpretation of definite descriptions is determinable through the principle of Relevance and, thus, the inferential process is encapsulated, as predicted by the modularity hypothesis.

Conclusion

It is not unusual to end a chapter by voicing further questions which the discussion have served to raise. Why is it that the human mind is programmed to store and use ‘default interpretations’?. I think this has to do with the number of simplification principles which the mind uses in order to reduce the complexity of the reality with which it is ordinarily confronted. Default interpretations can be seen as an attempt to model reality in a more simplified way. Contextual augmentations have the potential to calibrate inferences, to make them suited to particular contexts, to add richness to the schematic nature of basic pragmatic inference.

Chapter 5

Further reflections on Semantic Minimalism: Reply to Wedgwood

Abstract

In this chapter I discuss a paper by Wedgwood in which he considers the possibility that Relevance Theory and Semantic Minimalism share at least some common resources. I maintain that the two theories have different aims and different orientations and that it might be fruitful to understand why Cappelen and Lepore stick to Semantic Minimalism despite the various objections levelled to their theory. I explore certain minimalist solutions along the lines of considerations by Michel Seymour, adopting Jaszczolt's considerations on parsimony of levels of interpretation. I assume that logical forms contain certain variables which can be filled (or saturated) in context. As a final proposal, I broach the idea that pragmatic enrichment at the level of the predication can be avoided by resorting to a more complex enrichment at the level of the subject. I resort to ideas by Jaszczolt (specifically POL), to argue that parsimony considerations require that enrichments be operated at the level of subjects, if possible, thus avoiding a less parsimonious view according to which both subject and predicates should be enriched.

Introduction¹.

In this chapter, I reply to an important and also intellectually stimulating paper by Wedgwood (2007). Wedgwood's paper is of importance because it throws light on certain structural similarities between theories which belong to different conceptual domains, philosophical semantics and cognitive pragmatics. It creates undoubtedly important bridges between the two theoretical constructions and provides a very

¹ Abbreviations are explained in Appendix 1.

intelligent critique of Cappelen & Lepore's (2005) book. My difference with Wedgwood is one of point of view and of emphasis. Also, in some cases I point out what I take to be the weakest parts of his paper. I am sure he will find a way to refine those points and to answer them from his point of view, which is different from mine.

The three main threads of the chapter are as follows:

I disagree with Wedgwood's conclusion that there is more common ground between C & L and RT than there initially appears to be;

I disagree with Wedgwood's contention that RT's semantics is more minimal than C & L.

I hold that C & L's tests for context-sensitivity are robust, and thus that 'ready' and certain other apparently context-sensitive expressions have an abstract, context-insensitive semantics, which needs to be coupled with null prepositional phrases, which, in turn, require saturation.

Semantic minimalism, as proposed by C & L (2005), has met vigorous opposition on the part of Radical Contextualists. Semantic minimalism is the view that the semantic content of a sentence *S* is the content that all utterances of *S* share in virtue of what is contributed by their words and the syntactic relations among constituents². 'It is the content that **all** utterances of **S** express no matter how different their contexts of utterance are. It is also the content that can be grasped and reported by someone who is ignorant about the relevant characteristics of the context in which an utterance of *S* took place' (C & L 2005, 143). Semantic Minimalism claims that semantic content is pragmatically affected only with respect to standard indexicals,

² The addition of 'in virtue of what is contributed' serves to overcome the objection that – unless such a modification is added – one could hold that all utterances of a sentence *S* shared content *p* but that *p* was pragmatically enriched in ways not traceable to standard indexicals.

Wayne Davis (personal communication) considers that this definition must be conjoined with another definition which Cappelen & Lepore provide: Semantic Minimalism holds that "the semantic content of a sentence *S* is the proposition that all utterances of *S* express (when we adjust or keep stable the semantic values of the obvious context sensitive expressions)" (Cappelen & Lepore 2005: 2).

and that semantic content is propositional³. In his review of C & L, Recanati (2006) writes:

Like most philosophers of language, C & L are minimalists, but they defend a fairly radical version of minimalism, close to the literalist ideal. Instead of multiplying indexical variables, as other minimalists do, they maintain that there is only a short list of context-sensitive expressions and that sentences not including any of them are such that their meaning is their truth-conditional content (Recanati 2006: 21).

Recanati finds it strange that the authors should say that sentences such as ‘John is ready’ express complete propositions, while refusing to say what these propositions are (in detail), apart from saying things such as that e.g. ‘John is ready’ expresses the proposition that John is ready.

I agree that C & L should provide more than their mysterious story. At least one should try and provide some elucidations of the meaning of words such as ‘ready’.

The main argument against contextualism is that, once we fall into the trap of contextualism, the kind of argument used to show that the literal meaning of an expression is not sufficient to provide a full proposition will also show that the truth-conditional meaning arrived at through contextual augmentations will not be enough, as more stringent truth-conditions are required. Presumably this regress should be extended *ad libitum*. This is also another mysterious point. Even if we grant that the contextualist’s claims generalize too much (as the authors say), this does not amount to saying that an infinite regress is generated.⁴

³ One should note that the shared content assumption is not shared by all versions of minimalism (for instance, it is rejected in Borg (2004)).

⁴ One could cast doubt on my idea that there is a limit to what contextualist claims can say about meaning. Consider, for example, one of Cappelen & Lepore’s examples: ‘John went to the gym’. Did he go to the vicinity of the gym? Did he enter the gym? Or suppose we grant that the example is so heavily context dependent and that we resolve for this type of context-dependence, by saying that on a plausible interpretation the speaker meant that John entered the gym. Is not there a limit to what context-dependence could say about this example? Should we continue the context-dependent claims and say that underdetermination is there because we are still able to articulate the purpose of getting into the gym? Ok. Let us suppose that we also provide a purpose constituent: ‘John went to the gym to keep fit’. Let us suppose that we can go on providing implicit constituents until our imagination allows us to do so; is not there a limit to the number of constituents we can provide by free enrichment? My intuition is that it is not possible to add implicit constituents *ad libitum*, for one thing: the intentions of the speaker are finite and, thus, going on to add implicit constituents will not realistically model finite speaker’s intentions.

Despite various shortcomings, my view is that we should very carefully consider what C & L say and we should at least accept the various tests they propose which will enable us to detect context-sensitivity or, alternatively, context-invariance. This is not to say that the discussion of such tests should be uncritical and it is possible that such tests will be refined considerably in the years to come. In this chapter, I try to assess the validity of some criticism levelled to C & L on the part of Wedgwood (2007).

1. Some background.

In this section I shall mainly outline the main arguments of the paper by Wedgwood. The main aim is to create the context for the discussion of very specific and technical points later on in subsequent sections. In this way, it will not be possible to say that I have extrapolated excerpts from the general discussion by focusing on parts of the discussion which are not particularly crucial to the general line of the discussion.

Wedgwood sees RT and C & L's positions as being compatible. If this is right, Wedgwood takes this to be an important result as it represents the convergence of two major distinct approaches: a philosophical perspective on semantics; a cognitive perspective on pragmatics.

For Cappelen & Lepore, context sharing is the very foundation of communication. RT has a similar concern with content-sharing, even if this is fundamentally the result of the combination of encoded meaning and inferential processes.

In general, Wedgwood argues that C & L's outlook is essentially very similar to that of RT, despite C & L's persuasion that they differ substantially from RT.

According to Wedgwood, there is a difference between the two approaches. This is constituted by C & L's notion of minimal content. Nevertheless, Wedgwood argues that this is largely equivalent in function to RT's encoded meaning. For C & L, minimal content is propositional, while encoded meaning is not necessarily propositional. The key property of minimal semantic content is that it should be conveyed consistently across all contexts. According to Wedgwood, RT's encoded meaning does the same job.

Wedgwood agrees that RT is an example of what C & L calls radical contextualism in so far as it accepts that the influence of context on propositional meaning is pervasive and ubiquitous. According to RT, the intrinsic content of a

linguistic form is systematically sub-propositional (it must be complemented by pragmatic inferences to achieve full propositionality).

Wedgwood expatiates on the basic assumptions of RT, its character as a theory that studies ‘mind-reading’ processes in which a hearer must attempt to infer the speaker’s communicative intentions on the basis of knowing that the speaker relies on knowledge of the language and on propositions that are mutually manifest to both the speaker and the hearer.

Wedgwood stresses that RT considers that what is understood to be communicated by any utterance is heavily context-dependent and that, despite this, people succeed in communicating with each other without any special problem. This is ensured because communication is governed by some set of principles enabling people to recognize the intended meanings.

Wedgwood considers that despite the emphasis on context, RT admits that some meaning is encoded in linguistic forms. The meaning encoded in a given linguistic form is sometimes called its ‘logical form’ or, as Wedgwood prefers, ‘encoded meaning’.

According to Wedgwood, RT assumes that a speaker relies on assumptions which he takes to be mutually manifest in his expectation that the hearer will recover the intended meaning.

According to RT, interpretation proceeds according to a presumption of optimal relevance, where relevance is conceived of as a ratio of communicative rewards to processing efforts. These rewards are what RTs call ‘cognitive effects’ (such as strengthening or weakening/eliminating existing assumptions or providing new assumptions that interact with existing assumptions to allow new deductions).

Words and structures of a given language can be taken to convey a variety of things in different contexts, depending on the assumptions which are brought to bear on the interpretation process. This means that linguistically encoded forms constrain but do not (fully) determine what is conveyed on a given occasion.

According to Wedgwood, RT accepts that there are two distinct notions of content: the inputs and outputs of the communication process, that is to say encoded and inferred meanings.

These two levels are sufficiently similar to C & L’s position, which they call semantic minimalism which combines with speech act pluralism, the idea that the same sentence can be used to make a variety of (different) speech acts in different

contexts. C & L too recognize two important levels of meaning: minimal semantic content and speech act content.

A point of divergence between C & L and RT is that the former take content sharing to be guaranteed by the conventional means of language. RT, instead, does not agree that communication involves the sharing of (identical) thoughts, but only similarity of thoughts. Wedgwood takes RT to be a more realistic and psychologically plausible theory as it reflects our daily experience quite accurately. In RT the goal of communication is to enlarge the interlocutor's mutual cognitive environment; that is, communication is successful if two people can tell that they have some more assumptions in common than they had before. This does not require the exact reproduction of thoughts.

2. Discussion of Wedgwood's response to C & L: general considerations on semantic minimalism.

The paper gives the impression that Wedgwood is more favourable to RT than to C & L while he attempts to establish the basic tenets on which the semantics/pragmatics debate ought to rest. In the attempt to reconcile the two theories, Wedgwood claims that moderate contextualism is opposed to both Semantic Minimalism and RT. It is doubtful that by creating a common enemy one can reconcile two radically different theories. For one thing, moderate contextualism is opposed to Semantic Minimalism only in so far (as proven by C & L) one slides from there into Radical Contextualism. It might be legitimate to hold the view that C & L are against all forms of contextualism and, in particular, they go to some length to distance themselves from moderate contextualists. Even if MC did not always slide into RC, MC already goes beyond their belief that one should limit contextual influence to basic indexicals (or anyway, to elements of the Basic Set of context-sensitive expressions). However, given that C & L accept pragmatic intrusion at the level of pronominals and deictic elements, in principle they should accept moderate forms of contextualism provided that they are reducible to usage of pronominals⁵. It

⁵ Someone might say that this is something that C & L explicitly reject. Yet, as my quotation of Seymour makes clear, it is reasonable to suppose that Cappelen & Lepore are compelled to accept that when linguistic structure mandates a gap, such a gap must be filled by saturation. Of course, it is possible to have different views of the various examples discussed here or by Stanley. However, in principle Cappelen

is the sliding from MC to RC that crucially scares them. This assertion should be properly evaluated in the background of important considerations by Seymour (2010):

In more controversial cases like “Mary has had enough” and “John is ready,” it could be argued that their incomplete character may be explained by the implicit presence of a demonstrative expression like ‘this’, so that the sentences should read “Mary has had enough of *this*” and “John is ready for *this*.” Here I follow Capone’s minimalist explanation of the nature of such incomplete propositions. (Capone 2008). If they could be interpreted as implicitly containing empty slots that can be interpreted as demonstratives or discourse-deictic anaphoric expressions, sentences expressing incomplete propositions would indeed be harmless for minimalism. The reason is that they would be analysed as implicitly involving expressions belonging to the basic set. “John is tall” and “Mary is rich” would also contain implicit semantic empty slots calling for completion by a particular reference class. These sentences should perhaps be analyzed as “John is tall (relatively to *this* reference class),” and “Mary is rich (relatively to *this* reference class).” Quantified statements like “All came for breakfast” would be an elliptical form for “All of *them* came for breakfast” (or “*They* all came for breakfast”), and would thus also be implicitly containing expressions belonging to the basic set of indexical expressions. (Seymour 2010: 2872).

Now, it appears to me that the cases above are reducible to semantic items having invariant semantics combinable with implicit (or null) elements that appear to belong to the Basic Set as formulated by C & L. We do not need to extend the basic set and the assumption that sometimes expressions of the basic set are null from a phonetic point of view is needed independently of the specific issue we are confronted with here. Nobody disputes that we have null elements such as PRO or pro in English or in Italian. But even in abstraction from grammatically stipulated null categories, we have free recourse to null categories as in ‘Look’ (look at that) or as in implicit

& Lepore cannot deny that, when there is a gap, it must be saturated; in the same way they cannot deny that a deictic requires saturation; a pronominal requires saturation: pro and PRO require semantic values, either assigned by the grammar or by the linguistic co-text (Rizzi 1982).

arguments (The ship was sunk to collect the insurance). We may use a technical term for such null categories. One characteristic of these silent PPs is that they are elements present in semantic interpretation that are recoverable from context and without which semantic interpretation cannot arrive at a full proposition.

One might argue that my claim that we do not need to extend the Basic Set as formulated by C & L is confusing, as the implicit demonstratives that I am advocating to account for the context-sensitivity of ‘ready’, ‘enough’, comparative adjectives, etc. must be encoded in the meanings of these expressions (thus, I must be suggesting that the basic Set be extended to include these expressions).

Now we are at a fork. Either we say that, e.g., ‘ready’ is a context-sensitive expression, and extend the Basic Set, or we say, as I propose, that ‘John is ready’ has a meaning equivalent to ‘John is ready for that’.⁶ The semantics needs to say that ‘ready’ subcategorises for a PP (prepositional phrase), which can be explicit as in ‘John is ready for the exam’ or left inexplicit as in ‘John is ready’. Is this kind of uses rare in language? I would say not. I can reply ‘I am coming’ to my mother who says I should go to lunch. When a constituent is left out, it is reconstructed on the basis of context and of the grammatical patterns of the expressions which are actually present in the utterance (Thus I say ‘I am coming (to the kitchen)’ but not ‘I am coming (under the kitchen)’). Implicit arguments are not that rare in the language, as there are examples such as ‘The ship was sunk to collect the insurance’. (See Roeper 1987).

In other words, we need no special resources other than implicit arguments, that is null PPs, which are recognised as a grammatical category anyway. So we do not

⁶ Wayne Davis (personal communication) considers that we have two uses of ‘John is ready’. In one use, ‘John is ready’ is elliptical for ‘John is ready to go to dinner’. In this case, it is not plausible to think that ‘John is ready’ is elliptical for ‘John is ready for that’. Davis thinks that the sentence is elliptical for ‘I am ready for dinner’. He thinks we need to distinguish two views: (i) ‘ready’ is elliptical for phrases of the form ‘ready for NP’; (ii) ‘ready’ means ‘ready for that’. On (ii) ‘ready’ is indexical; on (i), it is not indexical and does not have a meaning although it is elliptical for phrases that do have a meaning. According to Davis, there would not be any independent meaning that ‘ready’ contributes to the meaning of ‘ready for dinner’ or ‘ready for that’.

My reply is that considerations of parsimony militate against Davis’s story. His story amounts to positing an ambiguity. At most we can grant an interpretative ambiguity. A common denominator could be ‘X is ready for x’ and allow that in some contexts x has the force of a demonstrative.

need to list, as the commentator implies, ‘ready’, ‘happy’, ‘enough’ etc. as part of the Basic Set. All we need to ensure is that null PPs are part of the Basic Set.

It could be argued that my proposal, in fact, amounts to accepting indexical elements at LF, in which case it would amount to accepting a position analogous to that put forward by Stanley. But SM is not a thesis according to which some contextual effects on propositional content are due to indexical elements at LF, but rather a thesis according to which propositional content is fully determined by the semantics of the overt material. Such contents are admittedly impoverished and skeletal, but they are the basics of a theory of meaning.

My reply to this objection is the following. Hidden indexicals as well as PRO and pro are present in the structure of the sentence. So there may be a minimalist story according to which some minimal propositions are calculable on their basis. There may be another story according to which the full proposition is fleshed out by assigning values to variables. This story which involves saturating variables is not in conflict with the minimalist account. Even assuming that minimalists can account for problematic adjectives like ‘ready’ in some other way, they would have to cope with the interpretation of pro and PRO. Such a story would involve variables at LF. What I think would characterize semantic minimalists as opposed to people like Stanley is the desire to keep truth-conditions minimal and to accept that things like ‘John is ready for that’ have minimal truth conditions even if no saturation has yet occurred.

3. RT and C & L on shared content.

Wedgwood goes on to write:

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

shown to be equivalent to that of RT in all crucial respects. As such, it is just as well for C & L that their ‘shared content’ arguments against RT fail, as any arguments of this kind that succeeded in undermining RT would also undermine their own position (Wedgwood 2007: 649).

The passage above is fairly problematic. It presents two points. Let us start with the first point.

Let us assume that this position can be summed up as follows:

The difference between RT and SM is that the former claims that propositional content ensues from inferential processes that take place in the hearer’s mind, whereas the latter maintains that such content ensues from meanings conventionally encoded in words and, hence, propositional content persists independently of particular instances of interpretation.

Now, it should be born in mind that an explicature is what is taken to express the propositional content and that, according to Carston (2002), at least, an explicature has two components: a semantically encoded component and a pragmatically derived component. Now, if things stand as Carston says, it is simply not true that propositional content ensues from inferential processes that take place in the hearer’s mind. At least part of this content is provided by the semantics. The difference between the two positions, as I see it, is that Semantic Minimalism stresses the conventional layer of meaning, while Relevance theorists stress the non-conventional layer of meaning. It should be conceded that in some cases at least propositional meaning is exhausted by the conventional meaning of a sentence; in some other cases, propositional meaning is not exhausted by conventional meaning.

There is a further complication in this picture. Carston is ready to accept that explicatures are developments of conventional meaning, even if she insists that usually conventional meaning under-determines semantic interpretation (and pragmatics must intervene to provide full propositional forms). However, there are cases in which one must decide whether the contribution of the semantics must be discarded or modulated – ironic utterances are cases in which literal meanings do not reach the level of propositionality. Cases of modulation à la Recanati (2010) also show that pragmatic interpretation may reach the level of compositionality.

So, it is not necessary that the conventional meanings of words will make up the linguistic component of the explicature. Both the conventional and the pragmatic

component of the explicature are subject to pre-semantic interpretation. However, it could be said that the conventional meaning of a sentence is where the pragmatic interpretation originated and that the human mind must be able to have access to the meanings from which the pragmatic interpretation originated. It is also plausible that the human mind, in constructing the propositional form, does not throw away those bits of language which served to arrive at the propositional form. To make an example, a hearer who encounters the utterance ‘The lion stood in the middle of the square’ will proceed from its literal meaning to its modulated meaning ‘The statue of the lion stood in the middle of the square’. However, the sentence ‘The lion stood in the middle of the square’ does not completely disappear from the hearer’s mind, since the effect of surprise crucially hinges on that linguistic construct. Poetic interpretation of metaphors, likewise, does not throw away bits of linguistic structure, as these contribute to poetic effects. Furthermore, a person who uses a sentence in its modulated sense and a hearer who understands it should be able to keep in mind not only the final product of interpretation, but the process as well and the linguistic sentence from which it originated. If this is the case, then one can argue that it is methodologically sound to keep the semantic and the pragmatic component of the explicature separate, while surely one must be able to merge them (at some stage).

Wedgwood, in his paper, characterizes RT as a theory of the hearer’s interpretation. It is not clear to me whether this is a point that is crucial to Wedgwood’s paper or not. It is possible that not much in the main argumentative line of that paper changed if this point was dropped. It is worthwhile discussing it nevertheless. Wedgwood’s insistence on the hearer’s perspective does more harm than good to RT. If I were challenged to provide arguments for this position, I should quote from Capone (2010a) on indirect reports. In this paper I argue that the process of interpreting indirect reports usually (albeit not always, as interpretation process must adapt to special contexts) prevents the speaker (the reporter) from using an NP (coextensive with the NP actually used in the original utterance) that would transform (and distort) the original speech act, as such a move would place excessive processing efforts on the hearer. Now, my understanding of indirect reports is that pragmatic principles constrain both the speaker’s behaviour and the hearer’s behaviour. Emphasis on the hearer would surely provide only half of the explanation required, perhaps the half which we need least, as after all it is legitimate to try and

explain the speaker's behaviour as well, and such a behaviour clearly incorporates expectations about the hearer's behaviour.

Wilson & Sperber (2004) write:

The central claim of RT is that the expectations of relevance raised by an utterance are precise and predictable enough to guide the hearer toward the speaker's meaning (Wilson & Sperber 2004: 607).

Since an utterance creates expectations of relevance, it is licit to deduce that the speaker's linguistic behaviour is what creates such expectations. Thus, both the speaker and the hearer must be taken into account in an explanation of how relevance works. After all, when the authors say that human cognition tends to be geared to the maximization of relevance, 'maximization' is best understood as maximization by both the speaker and the hearer. Someone might take issue with this, by saying that after all, when I say that human cognition is geared to the maximization of relevance, 'maximization' is best understood as maximization by both the speaker and the hearer, this is confused. Cognition does not involve a speaker and a hearer (only communication does). Despite my difference of point of view, I find this suggestion useful. Since I am obviously writing about cognition in the context of communication (I admit it is not my intention to write a paper about the hearer's cognitive powers in isolation from the speaker's cognitive powers), it is probably useful to distinguish the speaker's cognitive processes from the hearer's cognitive processes. However, I claim that cognitive processes follow a social path of intentionality, which means that the speaker in uttering a sentence is busy predicting what is going on to happen in the hearer's mind, while the hearer is busy understanding what the speaker meant in virtue of predicting the direction of the hearer's cognitive processes. It is correct to say that communication can be guaranteed by the fact that both the speaker and the hearer share the same psychological make-up (Recanati 2010). If Recanati's intuition must be respected, it is useful to see that cognition involves a mirroring relation between the speaker and the hearer.

If further evidence is needed, consider the following case, from Blommaert (2005). Blommaert considers a misunderstanding between himself and a research associate. They were at a conference together and they had rooms in the same hotel. They were speaking Dutch. Blommaert wanted to say that his associate had a nice balcony by a Dutch utterance equivalent in meaning to ‘You have a balcony too’. The word for ‘balcony’ in Dutch is ambiguous between ‘balcony’ and ‘breasts’. It happened that when he used this utterance, there was a woman walking in the opposite direction who was wearing a deeply cut summer top exposing parts of her breasts. Blommaert writes:

Unfortunately, my young female research associate had noticed the woman – she had picked up a contextualization clue – and the term ‘balcony’ suddenly acquired a very suggestive, sexually offensive, meaning, which called for extensive explanation and damage repair afterwards. My words had been placed in (or made to point to) a context which had altered their meaning, triggering a shift from a descriptive, neutral meaning for ‘balcony’ to an implicit, male sexist slang meaning (Blommaert 2005, 42).

This is clearly a situation where the speaker’s meanings (or interpretations) and the hearer’s meanings (or interpretations) diverge. This is a situation that perfectly models Wedgwood’s considerations on inferential work (on interpretation of utterances). Yet, I would like to claim that RT, like any other pragmatic theory, must not typically model situations in which there is a discrepancy between the hearer’s and the speaker’s interpretation. There are ways of expanding the discussion of Blommaert’s example that point to the fact that the addressee somehow goes wrong in interpreting intentions. In fact, she follows an individual path of interpretation preferring it to the social path of interpretation (see Jaszczolt 1999)⁷. Since Blommaert never pointed or looked at the woman with the deeply cut top, the hearer made reference to a contextual clue which was only (unilaterally) presumed. The context which she brought to bear on the utterance was her unilateral version of the context, as she had no evidence that the context was

⁷ Intentionality is a notion that makes sense of the content of speaker’s mind during the process of uttering a certain sentence as correlated to the utterance. Intentionality that follows the social path of interpretation is undispersed intentionality, according to Jaszczolt (2005). From Jaszczolt’s picture, we understand that the hearer is allowed to reconstruct the speaker’s intentions on the basis of conventions correlating sentences and meanings and conventions of use correlating utterances in context with interpretations. Speakers’ intentions and hearers’ reconstructed intentions should converge if both speaker and hearers follow the social path of intentionality.

shared. She also followed an individual path of interpretation by choosing the pejorative meaning, when she had a choice between a meaning which was in line with the intentions compatible with the public role assumed by the speaker (a professor speaking in the context of a conference) and a meaning which could implausibly relate to sexual matters. Presumably most Dutch people who use the word for ‘balcony’ are in Blommaert’s position; however, ordinary usage is never as ambiguous as Blommaert really wants to let us accept. In context, ambiguities are easily settled. One should always go for the social path of interpretation (as Jaszczolt 1999 says), as this is the only path of interpretation that can ensure that the speaker’s and the hearer’s meanings coincide. Since following the social path of interpretation involves the speaker’s duty to project meaning on the basis of conventions of language and of use and the hearer’s duty to interpret intentions on the basis of the same conventions of language and of use, convergence is guaranteed if both the speaker and the hearer abide by their duties. Presumably, there are also norms of interpretation besides the Gricean maxims, like the ones I tentatively put forward in Capone (2004a). The most basic ones of which are the following:

Show goodwill

Contextualise an action in such a way that it can be interpreted positively; if you do not find a context in which it can be interpreted positively, then at least allow for the possibility of finding a context in which the action can be interpreted positively.

Be constructive

Repair your coparticipant’s mistakes by attributing positive interpretations to her actions; in particular, adjust any interpretations of her actions by taking into consideration the intentions that can be plausibly attributed to her.

There is nothing ‘ad hoc’ about these norms for interpretations. The other day I happened to make a blatant mistake and I said ‘Hello’ to a woman in the street that looked like an acquaintance of mine. When I realized I had greeted the wrong person, I apologised. What followed was a mini interaction in which the woman intended to play down the event – categorizing it as a thing that can happen to everyone. If one wonders why such lengthy elaborations follow these incidents, the reason is clear. People want to be seen as showing goodwill – they will let errors pass and they will discard pejorative interpretations when it is possible. In the light of this,

what is striking about Blommaert's example is that it may serve to model a theory of language which is utterly different from RT – in fact, RT too needs a mechanism to ensure that in the standard case the speaker's and the hearer's interpretations coincide (or at least do not differ vastly). Wedgwood's picture of RT makes it appear suitable to modeling situations like Blommaert's, while, in fact, it is suitable to giving us a model of Blommaert's situation that tells us why the hearer went wrong, what was wrong in the interpretation process and why the speaker was perfectly entitled to receive an interpretation in line with his original intentions.

A possible objection to my approach based on good will is that a principle such as 'Show good will' is surely liable to change with context. In RT terms, the attitude of one's interlocutors is as contingent factor that may be mutually manifest or may be made so as a result of their behavior. Such attitude cannot therefore be constitutive of any general conversational principle.

My reply would be the following. Although I have not said anything about whether the principle 'Show good will' could be reduced to the principle of relevance (arguably it could be so reduced, given that such a principle would tend to maximize positive effects in the face of the same cognitive efforts; something to the same effect was argued for by Paglieri and Castelfranchi 2010 in connection with the maxim of Quality which could be subsumed, arguably, under the maxim of Good will), there are reasons to believe that 'Good will' is rooted in our culture and has roots in social practices. Just to give you an example, consider the example by Dummett, describing the evil teacher who asks a question expecting/desiring that the pupil will not answer it. In Capone (2010f), I argued that this a bad teacher, one who does not conform to the standards of benevolent teachers, who are the norm in society. He does not conform to the standard because he does not conform to the maxim of Good Will.

The second point by Wedgwood is more difficult. Wedgwood says that RT, despite what C & L claim, has an approach similar to C & L on content-sharing. Wedgwood accepts that the level of linguistic meaning due to the conventional meanings of words and grammar is what is common to both semantic minimalism and RT. Yet, the assumption that linguistic meaning is largely underdetermined makes the two theories sufficiently differentiated.

4. RT and C & L on the equation of semantic content with intuitions about speech act content.

Wedgwood (2007) discusses the (so called) Mistaken Assumption (by C & L):

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

What C & L apparently fail to recognize is that RT quite explicitly rejects this too (see further section 5.1). It may be that relevance theorists tend to express the idea in rather different terms – for example, ‘linguistically encoded semantics falls short of truth-conditionality’ – but since truth conditions are conventionally identified with what the speaker is taken to be committed to (i.e. to have asserted, claimed or stated) by uttering some sentence, this is in fact just another way of expressing the very same point (Wedgwood 2007: 654).

There are interpretative complications. What does ‘this’ refer to? Presumably it refers to The mistaken Assumption. ‘The same point’ also requires disambiguation: does the expression refer to ‘the same point as The mistaken Assumption’? Or does it refer to ‘RT quite explicitly rejects this too’? Even after such a disambiguation, there is very little sense in what is said. In any case, let us see if it is true that RT too rejects the equation of semantic content with speech act content. Carston (1991) says that a sentence such as:

(1) If the king of France dies and France becomes a republic, I will be happy, but if France becomes a republic and the king of France dies I will not be happy

has contradictory truth-conditional import and this defect is remedied by resorting to inferential pragmatics. In other words, truth-conditional import is calculated by trying to find out what kind of assertion the speaker made, what his intentions were in making this assertion. It is implicit in her view that the truth-conditional content of (the whole of) sentence (1) is not merely based on the truth-conditional content of

each of the conjuncts in each conditional, but there is a truth-conditional content at the level of the speech act that is (globally) pragmatically conveyed. Now, while surely both Carston and C & L agree that the semantics of ‘and’ is basically truth-conditional, it is not clear that for Carston the truth-conditional content of (the whole) of utterance (1) and the truth-conditional import based on sentential semantics do not come apart. Consider what Carston (1991) says about (1).

As Cohen (1971) has pointed out, if *and* is simply truth-functional and the temporal and causal connotations are captured by implicatures, then (22a) and (22b) should be contradictory at the level of explicit content (instantiating ‘If P then Q but if P then not Q’, and ‘Not P; P’ respectively) (...). However, these examples are not understood as contradictory or redundant. Those who wish to maintain an alleged implicature analysis have to say that the alleged temporal and causal implicatures contribute to the truth conditions of the utterance in which they occur (...).

So a clear difference emerges between, e.g. C & L and Carston. C & L would have to hold the view that there is a contradiction in (1) (which is then resolved at the level of the speech act). Carston, instead, holds the radical view that the pragmatics of conjunction is integrated into truth-conditional semantics. (I discuss this problem in Capone 2006).

I understand that Wedgwood would say, at this point, that Carston, like C & L, differentiates between sentential meaning and the level of the speech act. For Carston, the sentence would have to be contradictory at the level of the sentence, but not at the level of the speech act. But the question now arises, how is this possible? Do intentions have the magivcal powers to make a sentence that is contradictory non-contradictory? How can this be? It can be argued that the sentence is not contradictory in the first place. At this point, it is clear that Carston’s account and NOT Cappelen & Lepore’s account is confused. Cappelen & Lepore can argue that the sentence is not contradictory because it has temporal variables that can be filled and they are filled only at the level of the speech act. Carston, instead, projects the categories of the speech act to the level of the sentence. The problem does indeed arise from the habit of confusing speech act categories with sentential categories.

Now there is the further question of how we should use 'semantic content'. If we use semantic content in the way Wedgwood presumably does, this comes to (boils down to) semantic import. If we use it the way C & L do, this comes to propositional import. What Wedgwood commits himself to is a rhetorical fallacy (called 'equivocation'; see Stati 2002, 122): he attempts to defeat C & L by shifting the semantic import of their words. For RT 'semantic content' is linguistically encoded content, which may fall short of a proposition (a proposition being recovered only once rich pragmatic processing takes place). This distinguishes it from semantic minimalism which claims that linguistically encoded content will be propositional. Thus it is clear that 'semantic import' is understood differently by RT and by C & L, the former being committed by its use only to the underdetermined logical forms of semantics, while C & L allow a modest dose of pragmatic intrusion.

Of course, on both views, it should be granted that semantics should not be equated with the level of speech act content, since this invariably involves recourse to pragmatics. But the question whether a speech act can be associated, at least in some cases, with a minimal propositional content is undoubtedly important. The fact that we can communicate some messages in virtue of language in a non-miraculous way is of some concern for the theorist of language. The very notion of conventionality relies on the idea that we can express messages by using language. There should be some tough correlations between sentences and utterances – otherwise it would be absurd to create a sophisticated language system. Despite contextualists' attempts to place the onus of signification on contextual assumptions and pragmatic strategies, there is a reason why we do not have a magic word which in different contexts can express all the possible speech acts of the language. The fact that a language consists of (usually at least) fifty thousand words serves to illustrate that we pay a price for the possibility of expressing a potentially high quantity of meanings and that these meanings are structured on the basis of words – conventional tools at the basis of the semiotic process.

Words and syntax offered us an obvious way to escape the limitations of gestures, acts of pointing, pragmatic resources of all types, and to do so had to be based on the notion that language has at least some independence from contextual and pragmatic effects. This does not mean that language cannot be modulated in context, but that in principle it must be possible to interpret sentences on the basis of what they conventionally say in virtue of conventional meanings. This independence is what

guarantees that human beings follow a social path of interpretation when they proffer utterances. This independence is, furthermore, a model on which rules of language use are based. Pragmatics is not chaotic but essentially follows some rules or principles of language use and these are available to all users of the language and must be followed by all. It is only through the independence of language from pragmatics that language users could break out of the otherwise necessary lack of communication. The relative independence of semantics from pragmatics, furthermore, is what guarantees that pragmatics can have orderly effects on top of semantics, what guarantees that inputs and outputs are correlated through something which cannot alter completely the nature of the input. If it was possible to deform and freely alter the input, then nothing would prevent pragmatic effects from resulting into a chaotic magma in which the individual path of intentionality could prevail over the social path of intentionality, thus rendering communication impossible.⁸

5. On the equivalence of minimal semantics and encoded meaning in RT.

Wedgwood (2007) criticises C & L saying that, despite their attacks on RT, C & L's notion of semantic content is entirely equivalent to the notion of encoded meaning in RT. For the reader's convenience, I report an extract:

When C & L speak of truth conditions, they do not refer to the common (moderate contextualist) conception of 'what makes the proposition intuitively expressed by a given utterance true', but rather something that is indeed considerably more minimal. (1) is an example:

(1) 'Steel isn't strong enough' expresses the proposition that steel isn't strong enough (C & L 2005: 61).

Though this kind of statements appears less than dazzlingly enlightening, C & L insist that this is the level at which truth conditions should be stated, arguing

⁸ If am asked 'What did you mean by saying S in context C?', the set of my answers should be heavily restricted by the rules of language and the rules of use (the rules valid for the language game I am playing). I cannot reply by saying what might be convenient to me at that point of the interaction. The latitude in possible answers is severely restricted by the severe judgments of those who could reconstruct the meaning on the basis of what I said and the context in which I said that.

that it is the business of metaphysics, not semantics, to determine conditions on truth more precisely. The point for current purposes is that (1) shows how C & L's strict semantic argumentation – in common with their essentially communication-based 'shared content' arguments – is committed to the idea that the semantic content of a given sentence is just whatever remains common to its meaning across all contexts of utterance. There is no other way to interpret statements like (1).

By this definition, C & L's minimal semantic content is entirely equivalent to the notion of encoded meaning in RT. This is necessarily the case, since encoded meaning is by definition just that which is conveyed by a given linguistic form, irrespective of context. (Wedgwood 2007: 664).

According to Wedgwood, to say that some meaning is encoded is entirely neutral with respect to questions such as whether it has other properties, such as being propositional. All logical considerations militate against the equation of sentential levels of meaning with propositional meanings.

If we take what Wedgwood says seriously, we should concede that

(2) The sentence 'She is happy' ↔ It is true that Mary Thornton is happy⁹

The sentence 'She is happy' is equivalent to the proposition that Mary Thornton is happy. Presumably we have to assign the words 'entirely equivalent', as used by Wedgwood, a meaning which is radically different from what the words mean, say a speaker's meaning. Would this still be correct? I think it is not. There is no question that a sentence and a proposition can be equivalent. A sentence can serve to express a proposition in a given context, but, as pointed out by Strawson and by Frege (see Capone 2006), a sentence can be neither true nor false in itself, while a proposition is something that can be true or false. Sentences can be said to be equivalent in meaning only if they contain no indexical parts: alternatively we can say that two sentences are equivalent in meaning if, in case the indexical parts are saturated in exactly the same way by providing the same contextual value, uttering of one or the other sentence constitutes the same statement. The operation done by Wedgwood is

⁹ It was pointed out to me that there is no way to make this equivalence statement syntactically correct. But this is exactly the point of the example, there can be no equivalence statement of this sort.

very curious: the author subtracts important features from the programs of C & L and of Sperber & Wilson, and then says that by the phrase ‘semantic content’ they refer to more or less the same thing. The element subtracted from C & L’s program is partial, local intrusion, at the level of deictic terms; the one subtracted from Sperber & Wilson is radical intrusion, at the level of enrichment of predicate meaning. After these elements are subtracted, he shows that the programs are similar. What is not captured is the spirit of Sperber & Wilson’s program, which is to argue in favour of modulation, denying that verbs like, for example, ‘open’ can be assigned uniform truth-conditions. There is no denying that Sperber & Wilson’s program is more radical.

One could argue that I misrepresented Wedgwood’s point, which is that minimal content and encoded meaning are both assigned the role of being the content that is common across contexts, and that this means that SM will need to rely on similar processes as RT does when it comes to explaining how intuitive contents are arrived at in particular contexts. Wedgwood is not arguing that minimal content and encoded meanings are logically equivalent – but rather that the roles they are assigned by the two theories are structurally similar.

Now, acceptance of the above considerations, whether or not they represent Wedgwood’s stance, would surely be easier than accepting the text a fragment of which I quoted. Of course, one would wonder what role structural similarity would play in stressing the compatibility of the two theories being compared, if, as one should concede, full propositionality of the minimal level of meaning is an obstacle in the way of reconciling the two theories.

The types of arguments I presented above might be replied to by saying that it is true that RT sees the predicate of an utterance as generally underdetermining the property it is used to express, but RT denies that the ‘completion’ of the predicate contributes to semantic content. To justify my position I should argue why the minimal proposition should be considered the semantic content.

If semantic content is used as expressing the idea of the skeletal semantic representations of linguistic meaning, then I agree that Relevance Theorists do not claim that completion processes contribute to semantics. The challenge to justify the position that the minimal proposition should be the semantic content, is a challenge for C & L, as I am simply imagining how they might defend themselves.

We should start with making connections between semantics and the idea that semantics deals with communication and transmission of factual information. On this idea, two statements have the same content, not if they have roughly the same semantic meaning, but if they report the same situations, if they connect with the world in such a way that both statements are about the same situation. While surely we can use two different sentences to express the same situation, and we could also use the same sentence to express different situations, one point is clear: if two statements consist of the same sentences and they are proffered in the same context (say in replicas of the same context) and with the same intentions, the two sentences will be taken to make a report on a certain state of the world. If reporting on the world is the basic point of communication, we should be able to choose between a semantics which more directly connects with the transmission of information about the world and a semantics in which this connection is less direct. Suppose we choose an indirect connection between semantics and the purpose of transmitting factual knowledge. This means that semantics mainly deals with schematic information, which produces factual information only when coupled with contextual information. Is such a view plausible? And which view is more basic? The one in which semantics has a direct connection with the purpose of transmitting information, or the one in which semantics is only indirectly connected with the purpose of transmitting factual information? Choosing the indirect connection is a way of committing oneself to a view in which semantics is directly connected with some purpose, other than communicating factual information. But what purpose could this be? The alternative that language serves merely or primarily to construct relationships does not surmount these difficulties, as constructing relationships somehow relates to the expression of information about the world (about the way we feel about others).

And why should a looser relationship between semantics and communication of factual information be privileged? It is as if the theorist were to accept that language serves the purpose of providing information, but not information about the world, and that information about the world must be reconstructed on the basis of schematic, fragmentary information. But if this is the case, we must accept that this is the case even for prototypical utterances. A request for water as expressed by 'I am thirsty' provides or should provide bits of information, from which I reconstruct the puzzle of an intentional entity, the proposition that the speaker is thirsty.

The idea of a possible indirect link between language and communication makes it difficult to understand, for example, how the first intentional linguistic acts in the history of mankind were interpreted. If one starts with the idea of fragments of information being conveyed in semantic acts, one never really reconstructs the point of the utterance, as the utterance could be about ideas in general, rather than about facts. And yet, the intentional dimension of communication, as pointed out by Jaszczolt (1999) is a basis for all communication, a default principle of communication even. The very primitive acts of communication were interpreted as being about the world, and not about abstract ideas (and this made it possible for the hearers to utilize contextual clues for interpretation), but this could be possible only if there was perhaps an a priori principle, working in language (that is to say operative in the minds of language users) like the one produced by Jaszczolt, in favour of the intentionality of linguistic acts. It is true that in another book Jaszczolt (2009a) opts for merger representations in which linguistic and non-linguistic representations merge, to form a propositional dimension, but this is not to deny that prototypical, or primitive linguistic acts were devoid of intentionality or that in no cases the semantic representation corresponds to a proposition. I think it is clear that Jaszczolt is more interested in the complicated cases, but this is not like denying that there are simple cases.

6. A problem for semantic Minimalism: indexicals.

Wedgwood (2007) writes:

C & L's self-imposed commitment to the propositionality of minimal semantic content forces them to make this content less minimal than it could be (...). It also introduces a difference between their conception of minimal semantic content and RT's encoded meaning, as Carston (2007) stresses. In order to stick to what they recognize to be propositional meanings, C & L have to assume that minimal semantic content contains no unsaturated indexicals or ambiguities. The problem is that this immediately makes minimal semantic content something other than that which is shared across different contexts: the saturation of indexicals is a context-dependent process. (Wedgwood 2007: 666).

Let us see how this works in practice. Consider the following utterance:

(3) She is happy.

((3) is uttered by Mary who points to Ann when she says 'she').

According to C & L we could report Mary's utterance by saying (4):

(4) Mary said that she is happy.

Such reports can be uttered on the basis of the identity of the referent of the pronominal. Suppose this is a rule constraining indirect reports. The main question is whether there is identity of content between (3) and the reported part of (4), assuming identity of the referents of 'she' in the two utterances. If there is such an identity (and we suppose the answer is positive), then C & L's theory is a theory of abstract semantic content. The fact that semantic content is propositional is no impediment to having a view of content that is context-invariant. Propositionality simply means that we keep fixed certain elements which are tied to context while we use the tests for context-invariance. If, by keeping fixed the contextual contribution of pronominals, we obtain an invariant interpretation of the sentence, the results of this analysis will contribute to semantics and not to pragmatics. The fact that the semantic content of the sentence is propositional (in that it incorporates the semantic interpretation of indexicals and pronouns) is hardly a reason for thinking that C & L reject minimal semantics or for thinking that RT is a more minimal theory of semantic content. Since the contribution of context at the level of pronominals is factored out in the attempt to show that the reported utterance in an indirect report has the same content as in the original utterance, the theory has the level of abstraction required. Furthermore, it is simply not true that RT is a more minimal semantic theory simply because it factors out the contextual interpretations of pronominals, given the radical claim that the predication in an utterance is largely underdetermined (e.g. 'The car is red'). We do not grant that RT is more minimal at the level of pronominal interpretation, but even if we granted that, given that the predication act is largely underdetermined under RT, Wedgwood should be discouraged from arguing that RT is a more minimal semantic theory than C & L.

At this point the following objection should be answered:

But why couldn't RT just argue in parallel that the contextual contribution required by 'red' is 'factored out' – you get an invariant interpretation of 'that is red' so long as you fix both the referent of 'that' and the way in which the object satisfies 'red'. The point is that RT holds both pragmatic enrichments together, indexical saturation and predicate explicature). The question Wedgwood is pressing is why C & L treat the former as part of minimal semantics and the latter as part of pragmatic communication.

The answer to this question is very simple. C & L grant that pragmatic intrusion is indispensable when it comes to pronominals and similar expressions drawn from the 'Basic Set of genuinely context-sensitive expressions', while they do not grant that pragmatic intrusion occurs at the level of predicates such as, e.g., 'red'. Given the asymmetric status of pronominals and predicates, it does not come as a surprise that the pragmatics of pronominals is somehow incorporated into minimal semantics, while the pragmatics of predicates is relegated to the level of speech act interpretation. Presumably, the reasoning implicit in this asymmetric treatment is that pronominals and other deictic expressions do not contribute much to the thought expressed without pragmatic intervention, that is without relating the sentence to the context in which it is uttered. There is no thought at all, as Frege would say, without anchoring deictic elements to a given context. So, presumably in this conception there is an implicit equation between knowing the semantics of a sentence and knowing the thought it would express. A good question to ask is whether there could be genuine thoughts without anchoring deictic elements and pronominals to a context. The thought expressed by (5)

(5) She went to the cinema

cannot be the thought that a female person went to the cinema, since that would be expressible through (6) as well:

(6) A woman went to the cinema.

And (5) is different from (6).

A minimal semantics without provision of referents for pronominals would not describe thoughts, but would simply provide skeletal elements for the provision of thoughts. Instead, if my understanding of C & L is correct, they want to equate semantics with the vehicle for expressing thoughts, not just fragments of thoughts.

Wedgwood writes:

It is particularly strange that C & L should include the referents of overt indexicals in that meaning that is infallibly shared, since this is one of the most obvious sources of misinterpretations in language use (given an utterance of *He's ready*, surely a response like *Who? Mark or Paul?* is at least as likely as *Ready for what?*; yet, according to C & L's position, the former question should be pre-empted by 'shared content' while the latter is not). (Wedgwood 2007: 666).

The point is exactly the opposite of the one made above. C & L admit that there is a tightly restricted class of genuine context-sensitive expressions, such as pronominals, which have variable interpretation in relation to the actual context of use and linguistic expressions that are not context-sensitive. They deny that 'ready' is a context-sensitive expression.

Yet, one might insist that Wedgwood is exploiting C & L's claim that 'ready' is not context sensitive but that 'he' is. Given C & L's view, the referent of 'he' makes it into the shared semantic content while the determination of 'ready' doesn't, yet as Wedgwood points out hearers often don't know who is being referred to by an utterance of 'he' (just as they may not know 'ready for what') – but C & L's position should rule this out – if we grasp shared content for C & L we should know who the referent is.

There is no doubt that some hearers may be in a position not to know the referent of a pronominal. Suppose I am persuaded that behind me there is a picture of George Washington. It has been there for twenty years, and I expect it to be there while I am speaking. However, Mr Wedgwood to prove that C & L's theory is wrong suddenly removes the painting, while I say to my guests: 'He has a beautiful jacket'. When I turn round to point to the painting, I find out that George Washington is no

longer there. My referential intention, as Wedgwood would want to say, cannot be understood correctly by my interlocutor. Is this situation the norm? Is this the way we ordinarily use pronominals and deictic expressions? I would say it is not. It is possible that the addition of a semantic clause restricting the use of pronominals to accompanying demonstrative gestures is too strong. However, surely there is a norm preventing speakers from using pronominals unless the hearer can identify the intended referent of the pronominal through contextual clues that are unambiguous. Suppose I am a general reviewing an army. There are thousands of soldiers before me. Could I legitimately use: 'He is clever', unless there are contextual clues that enable the hearer to identify the referent? Presumably I could not.

So, while I agree that unlike predicates, pronominals do not have a content that can be grasped without recourse to a context, there must at least be some basic agreement with C & L that pronominals belong to shared meanings, in the sense that given the semantic resources available to speakers/hearers and knowledge of rich contextual clues, one is able to understand the meaning of linguistic expressions in context. Presumably to work perfectly, C & L 's view would have to be allied with a view of pronominals and indexical expressions in which semantics strictly regiments the content of a linguistic expression through a rule to the effect that pronominals and indexicals should be accompanied with unambiguous demonstrative gestures. Whether such a theory is desirable or terribly problematic is a topic for a different paper.

There is, obviously, the problem of how a context-sensitive expression can contribute a context-invariant meaning. But is this a real problem for semantics? While it is certainly true, as Wedgwood says, that the referent of a pronominal is provided by context, this is not to say that the semantics of pronominals cannot be furnished in a context-free manner. A very straightforward method is to let the pronominal refer to a thing x and to provide the semantics of the pronominal expression through a conditional:

(7) He is happy

is given the following semantics:

(8) If 'he' refers to x , then 'He is happy is true' just in case x is happy.

Following this reasoning, we can now understand how C & L can answer Wedgwood. The minimal semantics of a sentence includes the referents of pronominals and still remains context-invariant. To furnish the referent of a pronominal is just to furnish a specific value to 'x'. Replacing this value for x in (8), we obtain:

(9) If 'he' refers to John, then 'He is happy' is true just in case John is happy. 'He' refers to John.

The second clause of (9) is provided through pragmatics, but this does not destroy the context-invariance of the minimal semantics of (8), since this is perfectly represented in the first clause in (9). All C & L should do is use context-invariant semantics and couple it with a specification of the value of the variable involved in the logical form of the minimal semantics.

7. Indirect reports as a test.

Wedgwood discusses 'Indirect reports' as a way of testing intuitions about context-sensitivity. According to C & L indirect reports are a good test for context-sensitivity, as genuinely context-sensitive expressions, such as indexicals, as used in the original utterance, simply cannot be preserved in the indirect report. Here is an example:

(10) John: I am happy;

(11) John said that I was happy.

The sentence (11) cannot be a correct indirect report of (10), but something like (12) is required:

(12) John said that he was happy.

On the contrary, if a lexical item (say one in the predication) is not context-sensitive, it can figure both in the original utterance and in the indirect report, as the new

context of the indirect report does not change its interpretation. Here is an example from C & L (not a good one, actually):

- (13) John: Mary is ready.
- (14) John said that Mary was ready.

Furthermore, suppose that ‘ready’ was used in two (completely) different contexts, as in (15), (16), then it would still be correct to make the report as in (17):

- (15) John: Mary is ready (for the exam);
- (16) Ted: Mary is ready (to catch the train).
- (17) Both John and Ted said that Mary was ready.

This, however, should be understood as a quotational usage. The indirect reports promotes some common denominator interpretation, even if the two sentences were interpreted quite differently. I will return on the equivocation that allows C & L to obtain the wrong effects with the test in the case of ‘ready’.

However, Wedgwood (2007) writes:

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

There are various ways in which indirect reports of the kind that C & L rely on can fail to report a single proposition. For example, being the complement of ‘said’ does not force disambiguation of homonyms. I can truthfully utter the sentence ‘In C1 and C2, Nina said that John went to the

bank’ even if I know that Nina was referring in C1 to the financial bank and in C2 to the watercourse bank. (...) Significantly, one may easily do this with regard to C & L’s example also: In C1 and C2, Nina said that John is ready – but in C1 she meant that he is prepared for the exam, while in C2 she meant he has his raincoat on. (Wedgwood 2007: 669).

It is legitimate to raise such doubts. In fact, if indirect reports merely allowed quotational interpretations, then their efficacy would be destroyed. As Wedgwood says, they would no longer be efficient tests for propositionality. However, the very examples he uses persuade us that indirect reports like (18)

(18) In both C1 and C2, Nina said that John went to the bank

made on the basis of utterances such as (19), (20)

(19) Nina: John went to the bank₁ (financial institution)

(20) Nina: John went to the bank₂ (river course bank)

are not licit. I think everyone agrees that (18) should be different from (21):

(21) In both C1 and C2 Nina uttered the words: ‘John’ ‘went’ ‘to’ ‘the’ ‘bank’ in this order.

It would not even be licit to utter (22):

(22) In both C1 and C2 Nina uttered the sentence ‘John went to the bank’.

In fact, given the homonyms ‘bank₁’ and ‘bank₂’ there is not a single sentence, but there are two different sentences having distinct meanings.

Wedgwood (2007) also writes:

Similarly, C & L’s test fails to disambiguate use/mention ambiguities.

My intuitions are completely different. So called use/mention ambiguities are not really cases of ambiguity in the written language, where inverted commas are conventionally used to make reference to mentions of a lexical item. Oral language uses different ways of marking inverted commas (...). Given these facts, it would simply be absurd to use indirect reports in ways that ignore the use/mention distinction. (Wedgwood 2007, 669).

The attack on indirect reports as a test for a common meaning denominator is vaguely reminiscent of another attack on Cappelen & Lepore by Stanley (2005a). Stanley's position is misguided, because it focuses on a specific problem in C & L's approach and draws generalizations about a more general point. Stanley, after noting C & L's position according to which

The semantic content of a sentence *S* is the content of all utterances of share (p. 143)

moves on to ponder on the consequences of this statement. He writes:

This passage is rather unclear. Is it the content that all assertions of *S* express, no matter how different their contexts of utterance? If so, "every bottle is in the fridge" has no semantic content relative to any context, since there is no one proposition that is asserted by every utterance of the sentence (and certainly not, as we have seen, the proposition that every bottle in the universe is in the fridge, since this is never asserted). If the common content of all utterances of a certain sentence is not the content of any genuine speech act, what is the motivation for thinking that common contents are always genuine propositions, rather than just Recanati's "semantic schemata"? (Stanley 2005: 143).

However, this scepticism is simply defeated if one notices that C & L may well be wrong about quantifier domain restriction (a story like Stanley's being more palatable), while being right in the general claim that the semantic content of a sentence is the content that all utterances of *S* share (Perhaps the addition is required that we should confine ourselves to serious, rather than non serious uses).

To go back to Wedgwood, one of the merits of C & L is to say that we should extrapolate some common denominator meanings from the various and disparate uses of some linguistic (sentential) expression. It is essential for this virtue of the theory to survive that both Wedgwood's and Stanley's contentions be silenced.

We should now reflect on the meaning of 'ready'. Contextualists assume that 'ready' is more or less interpreted as 'ready for that'. This would explain why in a context 'John is ready' means that he is ready for the exam, while in another context it means that John is ready for the tennis match. I claim that contextualism along these lines does not seriously jeopardise C & L's work, as they, after all, admit that there is a restricted class of context-sensitive expressions, which they call 'The Basic Set'. Pronominals belong to this class. But are there ways to make sense of C & L's claim that 'ready' is not context-sensitive? It should at least be possible to discover that 'ready' is, semantically, not context-sensitive as contextualists say. I think it is useful to consult the Longman Dictionary of the English Language. The dictionary lists a number of senses, some of which appear to be more context-sensitive than the others. Let us start from the least context-sensitive:

- 'spontaneously prompt' (always has an answer);
- 'Notably dexterous or skilled' (he is very ready craftsman);
- 'Immediately available' (had little ready cash)
- 'prepared for immediate use' (dinner is ready')

Then we have some context-sensitive senses:

- 'forward or presumptuously eager' (he is very ready with his criticism)
- 'prepared mentally or physically for some experience or action'.

The real problem is to do with 'ready' what one should do with 'open', that is to say account for the various distinct uses of this word. Suppose we choose a sense such as 'prompt' ('disposed to act as occasion demands') as the basic sense from the others can be derived pragmatically. At this point, we can easily explain how Intercontextual Disquotational Indirect Reports make this basic sense emerge as a context-invariant one.

I should clarify that I am not saying that ‘John is ready’ is context-independent, but I am arguing that there is a meaning of ‘ready’ which allows C & L to obtain the wrong results from the test of indirect reports.

Someone might criticise my views suggesting that by treating ‘ready’ and ‘enough’ as having hidden indexical components, I am actually rejecting the results of C & L’s tests. This suggestion is too quick, and to some extent resembles Stanley’s tactic to generalise from a local problem to a general problem. It should be admitted in principle that even if there is a reliable testing procedure, the person effecting the test may be wrong in executing it. We should remember that tests should be used judiciously, as they depend, after all, on intuitive judgements. In the case of ‘ready’, C & L strain their tests of context invariance. When we say that Julia and Mary are both ready, given that Julia is ready for her English exam and Mary is ready for her maths exams, we are certainly right **on the interpretation** that they are both well prepared for their respective exams, but saying that they are both ready in the sense that they are both ready for *x*, *x* being the (specific) thing each intends to do is surely false. It may very well be that there is an equivocation of meanings in the way C & L use this test, but this does not mean that all users of the test should be led away by things such as equivocation (of meaning).¹⁰

8. Justification for Semantic Minimalism.

Last, I want to discuss Wedgwood’s following assertion:

Consider again the kinds of observation that C & L claim make shared context indispensable: co-ordinated action, collective deliberation, linguistic communication justifying beliefs, holding people responsible for what they

¹⁰ As Zwicky & Sadock (1975: 14) say, “If the semantic representation of certain sentences lack specification of some piece of meaning, then the applicability of transformation to them cannot possibly depend on whether or not this piece of meaning is present. If a sentence is unspecified with respect to some distinction, this lack of specification must be preserved by every transformational operation. But if a sentence is ambiguous, then it is possible for a transformation to apply in some but not all, of the cases, so that the effect of the transformation is to eliminate one or more understandings of the sentence”. In our case, the crossed reading can be obtained because the underspecified interpretations are preserved by the transformation of conjunction reduction. Since under-specification is preserved, the crossed readings can obtain.

say, and so on. Is minimal semantic content capable of supporting such things? If so, then an engineer who says ‘Steel is strong enough’ is held responsible just for that proposition (if we call it that) which is expressed by any and every utterance of this sentence. If the roof collapses, it should be considered reasonable of the engineer to point out that I just said steel is strong enough; I never said strong enough to support the roof. Such utterances would usually be described using words other than ‘reasonable’. (Wedgwood 2007: 670)

While in principle we could agree with Wedgwood that more than minimal semantics is required for language to stimulate coordinated action, two things ought to be noted here:

C & L might reply that minimal semantics is like the pillars and beams of a house; the house can’t exist without them. So communication is not possible without minimal semantics. Coordinated action as triggered by language implies a shared understanding of at least the basics of the utterances: that is to say their minimal semantics.

Second, in replying specifically to the point about the falling roof, an engineer who says ‘Steel is strong enough’ presumably means that ‘steel satisfies one’s needs. Then the pragmatics of his utterance will intervene in securing the interpretation that steel is strong enough for the roof. However, even if pragmatics did not intervene to secure such a specific interpretation, the engineer cannot be accused of incompetence or of having said something irrelevant.

9. On the compatibility between Radical Contextualism and Semantic Minimalism

Jaszczolt (forthcoming) claims that there is compatibility between Semantic Minimalism and Radical Contextualism. The compatibility in question mainly stems from the fact that (or as a consequence from the fact that) she abandons the syntactic direction principle, that is to say the assumption that the enriched propositions obtained through radical pragmatics should be considered developments of logical forms provided by sentential semantics. In other words, if I understand Jaszczolt well, in many cases, arriving at an enriched proposition is not a matter of taking a logical form as a point of departure and arriving at an enriched logical form, but it is

a matter of going beyond what the initial logical form says, of bypassing that logical form. Consider the examples

(23) You are not going to die

(24) Child: Can I go punting?

Mother: You are too small

It seems implausible that the hearer goes from the literal logical form ‘You are not going to die at t’ to the non-literal ‘You are not going to die from that cut’ or that in (24) the hearer first recovers ‘You are too small’, whatever its context-invariant meaning ought to be, and then moves on to the contextualised meaning ‘You are too small to go punting’.

There are also many other examples, from Capone (2009), where recovering an explicature is not just a question of developing a logical form, but is a question of bypassing it altogether:

(25) Sicily is a triangle

is one such case. In this case the enrichment is not of an augmentative type, but of a subtractive type, as it is implausible that Sicily is a literal triangle. So, it appears that the assumption that an explicature should be a development of a logical form in the sense that the logical form it starts with appears intact in the explicature must be abandoned.

I am aware that more than one approach will be available to this example, so I propose to use other examples, to show that the logical form of a sentence can be bypassed altogether in semantic interpretation. Cases of irony can prove that what is literally said need not be literally expressed and that, therefore, all cases in which a literal meaning is projected, heavy pragmatic processing is required.

The point by Jaszczolt that implicatures and explicatures can be calculated at the same time in the same merger representation is a bit more contentious. One of her best examples is the following:

(26) Everybody is going to Egypt this spring.

She claims that the main point of the assertion (the primary meaning) is that the interlocutors should consider going to Egypt this spring. She argues that primary meanings should be allocated a prominent place in a contextualist theory of meaning. That preoccupation, for Jaszczolt, seems to be incompatible with the syntactic direction principle. This point seems to me to be less well established, UNLESS we take her claim to be that a merger representation merges information deriving from explicatures and implicatures and that in many cases the literal interpretation is overridden by contextual considerations.

Let us return to the compatibility between semantic minimalism and radical contextualism. Once we agree with Jaszczolt that minimal semantics lacks psychological plausibility, a point that I think is conceded by C & L (2005), there are no problems in understanding how contextual considerations can override minimal semantics and how minimal semantics is absorbed and transformed into richer and often radically different propositional forms.

According to Jaszczolt, Semantic Minimalism is a project which is busy with truth-conditions at the sentential level, while radical pragmatics is busy with truth-conditions at the utterance level. In the same way in which semantic minimalism adheres to a compositionality principle, radical pragmatics adheres to a compositionality principle at a different level, given that the enriched propositions contain constituents which are not present at the level of logical form. It goes without saying that the compositionality principle employed in radical pragmatics is not a level of syntactic combination that belongs to surface structure (or to the logical form of the sentential semantics), but presupposes a level of syntactic combination that is (in principle) usable at the level of surface structure (or of the logical form of the sentential semantics).

According to Jaszczolt (2005), merger representations have the character of compositionality. This is to be accepted, as some implicit constituents, hence structure, is imposed on merger representations; we should suppose that the combination of certain semantic interpretations with certain pragmatically supplied constituents should follow compositionality. Consider the following example:

(27) That piano is better.

In (27) there is a missing constituent, and this one must combine with ‘That piano is better’. Now, in examples like (27) the kind of compositionality which the merger representation displays is exactly the same which the full sentence would exhibit. (There are, of course, infinite discussions on whether (27) is a complete sentence or not; what is incontrovertible is that the full alternative to (27) has the same compositional structure as an utterance of (27)). However, Jaszczolt claims that in some cases there are some mismatches between the sentence and the utterance, and that while compositionality cannot be held at the level of the sentence, it can be held at the level of the merger representation. The best examples of this mismatch are belief reports. Jaszczolt’s ideas derive from Frege. Compositionality, at the level of the sentence, in the case of (28) breaks down.

(28) John believes that Mary is happy.

It breaks down because, by the compositional picture, when one replaces ‘Mary’ with a coreferential NP, one obtains (or may obtain) a false statement. So, a picture in which we compositionally build up the meaning of the sentence/statement by replacing each constituent with the appropriate value it has in context breaks down because such a compositional picture would require the possibility of substitution of identicals in slots associated with, say, proper names. At the level of the merger representation, further structure is provided, that quickly shows how it comes about that substitution of identicals is not always licit in intensional contexts. (Implicit modes of presentations are provided and integrated into the semantics, say as appositives, following Capone (2008)). Belief reports are ideal candidates for the task of showing that compositionality obtains at the level of merger representations. I would add that ‘de se’ beliefs are even better for this purpose¹¹. Although my views on belief reports do not exactly coincide with Jaszczolt’s, we agree that the so called scope ambiguities do not resolve the problem of compositionality. In fact, assuming

¹¹ My paper on ‘de se’ attitudes (Capone 2010b) shows that ‘de se’ constructions are cases of intrusive constructions à la Levinson. ‘De se beliefs’ like ‘John remembers being in Oxford’ are beliefs about the self – they have first-personal readings which are truth-conditionally different from ‘de re’ readings. In my paper I argued that the mode of presentation ‘I’ is implicit in de se attitudes and furnished through pragmatic intrusion. Pragmatic Intrusion is also instantiated in the internal dimension of PRO in ‘de se’ constructions. See also Lewis 1979.

scope ambiguities, it is possible to differentiate the ‘de re’ and the ‘de dicto’ reading of a belief report

(‘de re’ reading: John is such that Mary believes of him that he is clever)

(‘de dicto’ reading: Mary believes that John is clever).

However, a problem to be noted is that the ‘de re’ reading entails the ‘de dicto’ reading, and thus scope ambiguities do not explain at all the problem of opacity (there would be some kind of circularity). It is best to accept, as Jaszczolt does, that compositionality is at the level of merger representation, where an implicit mode of presentation motivates opacity in the case where intentionality is dispersed (the strongest level of intentionality correlating with ‘de re’ readings, according to Jaszczolt). Does this amount to saying that the compositionality principle applies only at the level of merger representations? Should we reject the compositionality principle at the level of sentential semantics? (However, it should be born in mind that sentential semantics often consists of logical forms that do not match with the veneer of sentence structure). There is a sense in which the reply to this question is positive. If we do semantics in the merger representation, compositionality in the merger representation is all that is required (Jaszczolt 2005, p. 72). However, I also sympathize with more traditional semanticists, who claim that compositionality is an important characteristic of semantics. I can comfort them by saying that, after all, their semantics is done at a level of abstraction, thus it is methodologically correct to attribute compositionality to sentence structure, even if it has been abstracted away from the merger representation, for methodological reasons. The compositionality we find in the sentential semantics percolates to the semantics from the merger representation. However, since semantics is done, for methodological reasons, in abstraction from the merger representation, it appears as if the compositionality is there in the semantics. And in a sense, perhaps not the one by Jaszczolt, it is in the semantics.

10. A final note

In this section I will sketch a solution to some of the problems discussed, while I cannot go into the details of such a solution. It appears to me that semantic Minimalism seeks to establish an asymmetry between subject and predicate positions. Subject positions are invariably subject to pragmatic intrusion, and there is nothing one can do about it. Furthermore, pragmatic intrusion provides a constituent of the proposition, whether minimal or at the level of speech act theory. Without such a constituent there can be no minimal proposition, so Cappelen & Lepore cannot really provide minimal semantics without such constituents. Since subject positions are usually positions for reference assignment and reference assignment is pragmatic (having to resort to a number of contextual clues), there is no expectation that the subject position can provide a constituent of thought without referential resolution; and there is no expectation that there can be a minimal thought without referential resolution and the assignment of a constituent to the subject position. The predicate position, instead, has got a different status. It is true that a predicate is not immune to pragmatic enrichment; but it is not as dependent on pragmatic intrusion as the subject. Many cases have been provided to try to show that predicates cannot really furnish minimal truth-conditions. However, I am not persuaded that these cases are really against Cappelen and Lepore. Consider the case of ‘The lemon is yellow’. This may well require pragmatic intrusion, but only at the level of the subject. We may enrich the proposition up to ‘The lemon’s peel is yellow’. Since the pragmatic intrusion is required at the level of the subject, the predicate is not affected. Could we extend a similar treatment to ‘John is ready’? Presumably, we need to transfer the pragmatic intrusion from the predicate to the subject. One way to do so is through an implicit apposition constituent: ‘John [who must take the exam] is ready’. We could extend this position further by positing a null prepositional phrase as sub-categorized by ‘ready’. We thus obtain ‘John [who must take the exam] is ready [for it]’. And now we have obtained a considerable advantage. While I previously thought that ‘ready’ subcategorizes the constituent ‘for that’, which is fundamentally deictic, the prepositional phrase ‘for it’ is anaphoric. This means that the constituent [who must take the exam] is to all effects part of what is said, given that it is indispensable for anaphoric resolution. This could explain, presumably, why ‘Mary is ready’ cannot really mean Mary is ready for that, for that and for that. The

enrichment process is constrained and part of its constraints is that one should make sense of a sentence by enriching the subject first and one can enrich the predicate by anaphoric resolution through materials provided as apposition in the subject. Furthermore, one cannot say ‘Mary is ready and not ready’ meaning ‘Mary is ready for that and not ready for that’. This presumably follows from the fact that the prepositional phrase ‘For it’ is anaphoric and thus it would be a contradiction if the first occurrence of ‘it’ and the second occurrence of it both referred back to the same constituent. ‘Mary [who must take an exam] is both ready [for it] and not ready [for it].’

If I am right about these data, then we have bumped into the deep question of constraints to pragmatic enrichment and where they come from. But now we must ask the following questions. Do all predicates follow the same pattern as ‘ready’? Why is it that predicates can be enriched only through anaphoric resolution, while subjects can be enriched through provision of apposition constituents (that is to say in a really free way)? Why is it that indexicality cannot occur in predicates (unless use is made of explicit deictic pronominals), while anaphoric pronominals are tolerated? In reply to the first question, it appears that things are very much the same with ‘happy’. We cannot easily say ‘Mary was happy but unhappy’, meaning ‘Mary was happy about this but unhappy about that’, as we have the same anaphoric pattern as before: ‘Mary [who found her jewel] was happy about it but unhappy about it. It is also clear that somehow the anaphoric pronominal incorporates a reference to the event of finding Mary’s jewel, which compels the semanticist to incorporate the event in the apposition close through a device such as ‘Mary, who was such that there was an event of finding her jewel...’. Now I address the question of why pragmatic intrusion in the predicate appears to be subordinate to pragmatic intrusion in the subject and why it can occur through anaphora and not through deixis. A tentative answer has to do with a principle which has been brought to my attention by Jaszczolt (1999): Do not multiply levels of interpretation. This seems to correspond precisely with Cappelen and Lepore’s notion of semantic minimalism, in which intrusion is granted for subjects but not or only minimally for predicates. In other words, POL (Parsimony of levels of interpretation) compels us to minimize the loci of pragmatic intrusion and to utilize an already existing and necessary locus. Since the subject is a necessary locus of pragmatic intrusion, reference being necessarily a

pragmatic process and the subject being the locus for an NP that is also a referential position, pragmatic intrusion must occur primarily in the subject and if pragmatic intrusion occurs in a predicate, then this process must be subordinated syntactically to the subject position. Subordination allows us to keep the level of pragmatic intrusion in the predicate to a minimum (anaphoric resolution), while the pragmatic intrusion occurring in the subject is more radical.

Nelson Morales (personal communication) has thought of ways of falsifying my considerations. He thinks that loosening is such a case. He thinks that, in certain circumstances, a statement such as 'Mary is ready and not ready' is not contradictory (consider the interpretation 'Mary [who must take the exam] is [physically] ready but [emotionally] not ready'). His remarks are certainly useful, even if his example does not necessarily require loosening. As should be clear to anyone, contradiction is a property of statements and NOT of sentences. Someone who says 'Mary [who must take the exam] is [physically] ready and not [emotionally] ready [for it]' introduces some modifiers and this kind of modifier introduction is one of the various manifestations of pragmatic intrusion, not necessarily loosening. Notice that the anaphoric phrase [for it] stays there all the time and even in his example, one is prohibited from interpreting: Mary [who must take an exam] is physically ready for this [the exam] and emotionally not ready for that [getting married].

His example is very interesting and useful, as I said. Loosening applies to predicates like 'a triangle', when one says things such as 'Sicily is a triangle'. However, these cases are different because what is involved is a modification of the predicate to make the predicate fit the subject: Sicily has the rough shape of a triangle. Alternatively: Sicily's shape resembles a triangle. Strictly speaking, it would not be true to say that Sicily is a triangle. And again, one could contrive the case in such a way, that the pragmatic intrusion is at the level of the subject and not of the predicate. One, in fact, could have the following result:

Sicily's rough shape is a triangle.

This would very much be compatible with Jaszczolt's POL.

Let us now consider the following objection by an opponent:

The first point I would make is this: in the rest of the chapter you contrast indexicals (including pronominals) and predicates. However, in the final section you advert to a different notion of 'predicate'. Here 'predicate' identifies a sentence constituent ('THE predicate') in contrast to the sentence constituent you refer to as 'the subject'. On this latter use it corresponds to VP, assuming $S \rightarrow NP + VP$. It is this that leads you to refer to an asymmetry in C&L's treatment of subjects and predicates and, on that assumption, to your proposal.

But this (VP) sense of predicate is not the sense of 'predicate' appealed to in the rest of the chapter which I'll call "pred-1". Pred-1's can and do occur in the subject (and indeed all over the place, it covers common nouns, verbs, adjectives and adverbs). Any *descriptive* subject expression includes a pred-1. So I don't think C&L's position can be said to result in an asymmetry between subject and predicate - and that perhaps you need to look again at your proposal in the light of that.

Answer: it is true that Cappelen & Lepore speak of sub-constituent predicates, and of them they say that they are not to be modulated (their meanings are stable and are not in need of modulation). On the contrary for more radical treatments, predicates (at the sub-constituent level) require free enrichment because they cannot contribute to a full proposition; in other words, the contribution of the predicate to the proposition is sub-propositional. But then if my predicate roughly corresponds to a VP constituent, it is clear that the VP which contains as a sub-constituent a Predicate 1 also contributes a sub-propositional contribution; in other words, by applying that VP to the subject assuming that the subject can be used to identify a referent, one still obtains a sub-propositional element, that is to say an element which is not fully truth-evaluable. Since subjects are referential elements and reference always needs pragmatics means for being secured, it may be useful to transfer pragmatic intrusion from the level of the Predicate 1 and, subsequently, from the level of the next projection, VP to the subject, which, as I have already said, needs to secure its referent through pragmatic means.

Let us consider the next objection by an opponent:

I don't think you CAN "transfer the pragmatic intrusion to the subject" from the predicate, as you suggest, i.e. by interpolating a non restrictive (NR, appositive) relative clause. [1] Interpolating an NR relative does not actually enrich/modulate the referential semantics of the subject (unlike restrictive rel clauses); indeed, it has been argued that NR relative clauses aren't true sentence constituents anyway; [2] there is no limit to the range of the content of such NR clauses you might add (anything you might happen to know the referent of the subject); [3] as I understand it, you want to restrict the content of the relative clause in a way that might be relevant to the modulation of the predicate but that could be argued to be circular, either empirically or methodologically: first you have to identify the subject in order to know what modulation of the predicate might be relevant but you are arguing that the subject must be modulated in a way relevant to how the predicate is modulated.

Answer: why should we secure modulation in the subject by taking into account possible ways of enriching the predicate? The predicate does not require any enrichment if we are ready to enrich the subject in the proper way. When I said that 'Sicily is a triangle' may correspond to the proposition 'Sicily's rough shape is a triangle' I gave up completely the idea that the predicate or VP needed enrichment. I said that I transferred enrichment from the predicate to the subject as a way of speaking, as a way of contrasting my analysis with those of radical contextualists, but the transferral of the enrichment from predicate to subject, in fact, may never occur (otherwise, it would be correct to say that it would be circular). All we need is to enrich the subject.

The Obstinate Opponent argues that Non-restrictive relative clauses may not be relevant to pragmatic enrichment. Yet, if one makes use of them in the context of a sentence such as 'Mary is ready', which contextualists argue to be context-sensitive, and it can be demonstrated that they play a role in modulating the meanings of such sentences, it goes without saying that they must ancillary to pragmatic enrichment, whether or not, by themselves, they constitute cases of pragmatic intrusion.

Conclusion

I confined myself to the discussion of Wedgwood's points from the perspective of C & L. In general, we should be taught some salutary lessons by C & L concerning minimal semantics. This is not to say that we should abandon a contextualist approach to language. But perhaps it is convenient to make such an approach as compatible as possible with C & L's assumption that the Basic Set of context-sensitive expressions should be maximally restricted. There are some interesting context-sensitive expressions mentioned by C & L, such as 'enemy', 'friend', 'foreigner', 'outsider' etc. (C & L p. 1), but these are almost never discussed in the literature. I think we should concentrate on cases such as these or such as those dealt with in Capone (2006), and Capone (2008a).

The ideas on the semantics/pragmatics debate are quite fluid for the time being, but I have a propensity to follow Jaszczolt in her radical pragmatic perspective on language, trying to make her picture compatible with more classical theories on semantics. The idea that semantic minimalism and contextualism should not be enemies has recently been voiced by Recanati (2010). Of course, to enhance this compatibility one would have to accept that at least in many cases, but not in all cases, the meanings of sentences are truth-conditional and that at least in basic cases knowledge of the language will allow one to say what must be true for an utterance of a given sentence to be true. Contextualism, in such a hybrid theory, would be considered as a tool for enhancing the semantic potential of a language, rather than for arguing that language in itself is value-less. I remember distinctly Higginbotham's (personal communication) words in this respect – pragmatics is a **path** that makes truth-conditional semantics work. It is very useful to think of semantics and pragmatics in this way: as a path and a vehicle. You cannot use your car unless you have a road; the road is what makes your car useful. The road without the car is useless. Nevertheless, you can use your car in as many roads as you want to. This invariably involves a certain degree of autonomy.

Appendix 1:

Abbreviations

In this chapter, I make use of the following abbreviations:

SM = Semantic Minimalism

RT = RT

RC = Radical Contextualism

MC = Moderate Contextualism

C & L = Cappelen and Lepore

Appendix 2.

Wedgwood, in replying to my ideas, was bothered by the fact that Cappelen & Lepore allow pragmatic intrusion for indexicals but not for predicates. However, on my view that at least in some cases, predicates subcategorize for null prepositional phrases a limited form of pragmatic intrusion could be allowed for predicates as well. On my view, the asymmetry could be dissolved. Of course, this approach may be suitable for some cases; I am not saying that it should be suitable for all predicates. In fact, it is plausible that some forms of modulation à la Recanati should be accepted.

Chapter 6

Belief reports and pragmatic intrusion.

Abstract

In this chapter, I explore Bach's idea (Bach, 2000) that null appositives, intended as expanded qua-clauses, can resolve the puzzles of belief reports. These puzzles are crucial in understanding the semantics and pragmatics of belief reports and are presented in a section. I propose that Bach's strategy is not only a way of dealing with puzzles, but also an ideal way of dealing with belief reports. I argue that even simple unproblematic cases of belief reports are cases of pragmatic intrusion, involving null appositives, or to use the words of Bach, 'qua-clauses'. The main difference between my pragmatic approach and the one by Salmon (1986) is that this author uses the notion of conversational implicature, whereas I use the notion of pragmatic intrusion and explicature. From my point of view, statements such as "John believes that Cicero is clever" and "John believes that Tully is clever" have got distinct truth-values. In other words, I claim that belief reports in the default case illuminate the hearer on the mental life of the believer, that includes specific modes of presentation of the referents talked about. Furthermore, while in the other pragmatic approaches, it is mysterious how a mode of presentation is assumed to be the main filter of the believer's mental life, here I provide an explanatory account in terms of relevance, cognitive effects, and processing efforts. The most important part of the chapter is devoted to showing that null appositives are required, in the case of belief reports, to explain certain anaphoric effects, which would otherwise be mysterious. My examples show that null appositives are not necessitated at logical form, but only at the level of the explicature, in line with the standard assumptions by Carston and Recanati on pragmatic intrusion. I develop a potentially useful analysis of belief reports by exploiting syntactic and semantic considerations on presuppositional clitics in Romance.

Introduction.

As Mey (2001) says:

Pragmatics admonishes the linguistic scientists that they should take the users of language more seriously, as they, after all, provide the bread and butter of linguistic theorizing (...) (Mey 2001, 289).

It is in the spirit of this view that I write a chapter on belief reports and pragmatic intrusion¹. In this chapter, I discuss the issue of belief reports and propose to integrate it with the recent idea (mainly propounded by relevance theorists such as Carston 2002 and Sperber and Wilson 2002, but also, in different form, by Bach 1994, Levinson 2000, and Mey 2001) that the proposition expressed by an utterance (in a context C) is ultimately fleshed out (supplied on the basis of a skeletal semantic template, to use words by Carston) by recourse to pragmatics, which constructs missing constituents or expands the bare semantics of a sentence to resolve potential inconsistencies or absurdities. In particular, I propose that belief reports are cases of “intrusive constructions” (to use a term by Levinson 2000, 213), in that the truth-conditions of the whole depend on a pragmatic process of interpretation. In this chapter, I accept Sperber & Wilson’s view that as the gap between sentence meaning and speaker’s meaning widens, it increasingly brings into question a basic assumption of much philosophy of language, that the semantics of sentences provides straightforward, direct access to the structure of human thoughts (Sperber & Wilson 2004).

First of all, I discuss the hidden-indexical theory of belief reports by Schiffer (1995). Then I consider certain problems raised by Schiffer (2000) and Recanati (1993). I argue that these theories can be improved and that pragmatic intrusion can resolve the puzzles raised by the semantics of belief reports (see also Jaszczolt 2005,

¹ I was told that the term ‘pragmatic intrusion’ has negative connotations in the context of a theory in which the meanings of sentences are in general underdetermined. I do not take the view that a language, in order to be perfect, must match logical forms and propositions and thus I do not take intrusion as a sign of imperfection. It is true that that better terminology could be used, e.g. pragmatic inserts; but I am persuaded that the use of novel terminology may confuse readers who are used to books like Levinson (2000) or Carston (2002). So, I hope to be allowed to retain the term ‘intrusion’, imperfect though it may be.

126). I offer some views of pragmatic intrusion (mainly Carston's 1999 and Sperber & Wilson's 2002 views). Then I work out the details of the pragmatic analysis of belief reports. In particular, I offer an alternative to Salmon's (1986) view that conversational implicatures can explain away apparent cases of lack of substitutivity in belief reports, proposing that a communicative act triggers the expectation that it will be maximally relevant in that it will produce a high level of positive effects worth the hearer's processing efforts, compatibly with the speaker's preferences and abilities (Sperber & Wilson 1995; Blakemore 2000). It appears that the relevance theory approach best handles the cases of belief sentences used in thought (rather than in assertions), while Salmon's treatment, as Green (1998) points out, makes use of a Gricean explanation in terms of 'assertions' in conversational settings. Furthermore, unlike Devitt (1996), I propose that belief reports essentially illuminate the hearer on the cognitive state of the believer. The most important part of the chapter is devoted to showing that the presence of null appositives to NPs (within belief sentences) is required at the level of the explicature of a belief report. I mainly investigate cases of anaphora, control, and ellipsis and I extend the discussion to modal verbs such as 'ought' and to reflexive belief.

1. The hidden-indexical theory of belief reports.

Schiffer (1995) presents his hidden-indexical theory of belief reports. According to him, a sentence such as (1)

(1) A believes that S

expresses a three-place relation $B(x, p, m)$, holding among a believer x , a mode-of-presentation-less proposition p and a mode of presentation m under which x believes p . According to Schiffer (1995),

it is possible for x to believe p under one mode of presentation m while believing not- p under a second mode of presentation m' , and while suspending judgement altogether under a third mode of presentation m'' (Schiffer 1995, 248).

As Schiffer says:

This propositional mode of presentation is determined by modes of presentation of the objects and properties the proposition is about (Schiffer 1995, 249).

Schiffer admits the difficulties inherent in explaining what a mode of presentation (of a proposition) is and suggests that modes of presentations are mental representations that play certain functional roles.

Schiffer writes:

This theory is aptly called the *hidden-indexical* theory because the reference to the mode-of-presentation type is not carried by any expression in the belief ascription. In this sense, it is like the reference to a place at which it is raining which occurs in an utterance of “It’s raining”. And the theory is aptly called the hidden-indexical theory, because the mode-of-presentation type to which reference is made in the utterance of a belief sentence can vary from one utterance of the sentence to another (Schiffer 1995, 250).

Schiffer acknowledges that the reference to a mode of presentation is similar to the reference to a place in an utterance of (8):

(2) It is raining.

Endless discussions have been made on whether the place is or is not an implicit argument of the verb ‘rain’ in sentences such as (2), but I think that Carston’s view that ‘rain’ is not associated with a location variable is reasonable (being more parsimonious; see also Cappelen & Lepore 2006 on unarticulated constituents) The parallel between sentences such as (2) and sentences such (3)

(3) John believes that Mary is pretty

induces us to think that the mode of presentation associated with the embedded proposition of (3) is furnished through pragmatics (specifically through pragmatic

intrusion and explicature). Schiffer's statement that the mode of presentation type varies from context to context also points to its pragmatic nature.

2. A problem in the hidden-indexical theory.

Schiffer (2000) discusses a major problem for the hidden-indexical theory. The hidden-indexical theory entails that believing is a **three-place relation**, $B(x, p, m)$, holding among a believer x , a proposition p , and a mode of presentation **m** under which x believes p . The ordinary-language way of representing this open sentence is evidently 'x believes p under (mode of presentation) **m**'.

(...) and from this it follows that a singular term replacing 'm' in a true substitution instance of (3) [Ralph believes that Fido is a dog under mode of presentation m] would be the specification of an argument of the three-place relation expressed by the open sentence. (...) The intuitive point is that (3) clearly does not *look like* it contains a three-place verb with the specification of a third argument. Rather, it looks to be on all fours with a sentence like

(4) Louise hit Ralph under the influence of crack.

which is paradigmatically a sentence in which the singular term 'the influence of crack' is not the specification of the third argument in an instance of the three-place hitting relation (...) but rather merely part of the adverbial phrase 'under the influence of crack' (Schiffer 2000, 19).

Schiffer uses a Chomskyan diagnosis to check that 'under mode of presentation m' is an adverbial, and not an argument of the verb (believe). Consider the sentence (4):

(4) Mary gave the house to her husband.

We know that 'to her husband' is an argument of the verb 'give' because one can give the answer 'to her husband' to the question (5):

(5) To whom did you wonder whether Mary gave the house?

Now, Schiffer thinks that the answer “Under mode of presentation m” to the question
(6)

(6) Under what mode of presentation did you wonder whether Ralph believes that Fido is a dog?

does not show that ‘under mode mode of presentation m’ is an argument of ‘believe’ but shows that it is an adjunct of ‘wonder’. It would, in fact, be an elliptical answer for:

(7) I wonder under mode of presentation m whether Ralph believes that Fido is a dog.

Schiffer is aware that some scholars have disputed this syntactic test, yet he still maintains that there is a problem for those holding the hidden-indexical view.

Schiffer hopes to find a way out of this problem by following Recanati’s (1993) proposal. Recanati thinks that ‘believe’ expresses a two-place relation between a believer and a quasi-singular proposition. The quasi singular proposition is something along the following lines:

$\langle\langle m', \text{Fido} \rangle, \langle m'', \text{doghood} \rangle\rangle$

As we can see, the directly-referential singular terms of this quasi-singular proposition are **each** associated with a mode of presentation. However, modes of presentation do not contribute to truth-conditions. Schiffer modifies Recanati’s proposal by introducing the requirement that the modes of presentation m’ and m’’ be not specific modes of presentation but types of mode of presentation. Thus, where T ranges over modes of presentation types:

$B(x, \langle\langle T, \text{Fido} \rangle, \langle T', \text{doghood} \rangle\rangle)$ iff

$\exists m \exists m' (T m \ \& \ T' m' \ \& \ B(x, \langle\langle m, \text{Fido} \rangle, \langle m', \text{doghood} \rangle\rangle))$.

In other words, the treatment above makes it clear that the modes of presentation associated with ‘Fido’ and ‘doghood’ are not specific modes of presentation, but types of modes of presentation.

According to Schiffer the treatment above violates Recanati’s availability hypothesis: ordinary speakers ought to have some awareness of referring to such modes of presentation types. He writes:

Yet, I submit, it seems clear that the belief reporter has no such awareness. If asked what she was referring to in her use of ‘Fido’, she would not give any restatement that indicated that she was referring to a mode-of-presentation type (Schiffer 2000, 29).

As Schiffer is aware, this is not just a problem for Recanati’s analysis, but for the hidden-indexical theory as well.

The most crucial problem, though, is the following: where do the modes of presentation come from in an analysis such as:

<<m’, Fido>, <m’’, doghood>>?

Recanati and Schiffer appear to believe that the modes of presentation are in the that-clause complement of ‘believe’ independently of the semantics of the verb ‘believe’. Yet, we all agree that verbs such as ‘believe’ or ‘know’ create opacity, a problem which the hidden-indexical theory was contrived to deal with in the first place. In a sense, they are making it appear that it is the semantics of the that-clause that is responsible for a structure such as

<<m’, Fido>, <m’’, doghood>>.

Yet, they ought to be aware of examples such as (8)

(8)The judge decided that John Rigotti should die.

It seems reasonable to suppose that in this example the judge’s decision applies to the referent of ‘John Rigotti’ under any mode of presentation whatsoever. Yet, as

Williamson (2006) notes, it must be granted that if John Rigotti = The XYZ, then when John Rigotti dies so does the XYZ. But it does not follow that the judge decided that the XYZ should die. This shows that that-clauses, on their own, cannot be held responsible for opacity effects. The opacity comes from the verb e.g. ‘decided’. It seems to me that a more fruitful tack is to suggest that modes of presentation or modes of presentation types *are* built up locally within the NPs and VPs or APs contained in a that-clause embedded in a belief verb as **pragmatic increments**. They subsequently interact with semantic aspects of the verb of the main clause to create opacity effects. They also interact locally to incorporate aspects of the NPs, VPs, and APs of the that-clause. In fact, normally opacity effects are explained away under the assumption that the modes of presentation within the that-clause incorporate certain names. The proposal I articulate in this chapter assumes that verbs like ‘believe’ (attitudinatives, to use a term by Green 1998)) are semantically univocal and not ambiguous. It appears to me that by placing the burden of providing modes of presentation on pragmatics we can abide by Modified Occam’s Razor, which advises us not to multiply senses without necessity.

Salmon must be quite right in saying that modes of presentation are part of the pragmatics of belief reports:

(...) there is an established practice of using belief attributions to convey not only the proposition agreed to (which is specified by the belief attribution) but also the way the subject of the attribution takes the proposition in agreeing to it (which is no part of the semantic content of the belief attribution) (Salmon 1990, 233-234).

3. Pragmatic intrusion

Many authors have dealt with the semantics/pragmatics debate, but here I shall mainly expose ideas by Carston (1999) and Wilson & Sperber (2002). The reader will find an overview of other theories in Capone (2006). It is a pity that I cannot deal adequately with the views of Bach (1994), Levinson (2000) (see my review), Recanati (2004) (see my review) and of Stainton (2004a;b) which are also important.

Carston opposes the view that decoding utterances is merely a matter of coupling logical forms with pragmatic information:

A different view of pragmatic inference was suggested in the previous section, according to which this sort of inferential activity is an automatic response of receivers of (attention pre-empting) ostensive stimuli; it is but a particular instance of our general propensity to interpret human behaviour in terms of the mental states (beliefs, desires, intentions) of the behavior, which, in turn, is to be located within a bigger picture of general relevance-seeking information processing. According to this view, pragmatic inference is fundamental and the employment of a code (linguistic system) as an ostensive stimulus is a useful addition; it would not be reasonable to expect, nor would it be particularly desirable, that the forms supplied by the code should be eternal or even fully propositional (Carston 1999, 106).

No doubt, pragmatic information is useful in constructing fully truth-evaluable propositional forms, but compositional semantics plays a crucial role in the interpretation of sentences/utterances combining lexical with grammatical information. Even if Carston's claims are correct, and human communication reserves a large role for pragmatic interpretation in the construction of propositional forms, it is **methodologically** important to stress the role played by linguistic semantics and, in particular, by the lexicon and syntactico/semantic compositionality (Capone, forthcoming; Stanley, 2005). Carston tempers her view somewhat on p. 114:

The semantics/pragmatics interface is a representational level described as logical form or the linguistic semantic representation; it is standardly not fully propositional but rather a schema for the construction of fully propositional representations. Exactly what this looks like is, of course, an important question and not one that can be answered with any great conviction. A reasonable construal is of a structured string of concepts, configured along the lines of Chomskyan LF, perhaps indicating relative scope of quantifiers and negation, and with open slots marking constituents that must be contextually filled, as in the case of indexicals, quantifier domain, and many other elements (Carston 1991, 114).

Below are some of the examples relevance theorists use in support of the case for pragmatic intrusion:

- (9) The steak is raw;
- (10) Holland is flat;
- (11) Jane is a bulldozer;
- (12) He took off his boots and got into bed;
- (13) Writing an essay will take time;
- (14) Everyone went to the party.

(9) is true even if the steak is not completely raw, but only partially cooked; (10) is true even if Holland's surface is not, strictly speaking, completely flat but is flatter than most other European countries; (11) is obviously false, and thus a metaphoric interpretation must be accessed to consider it true; (12) is true in case the action of taking the boots off precedes the action of going into bed; (13) does not express the trivial proposition that writing an essay takes some time, but that it takes a considerable amount of time (and attention); (14) obviously does not mean that all human beings went to the party, but that all members of a certain domain went to the party; thus the domain of the quantifier must be suitably restricted by means of contextual knowledge.

4. Puzzles arising from belief reports.

If we abandon the hidden-indexical theory, there is no easy way to handle the puzzles of belief reports. I start to discuss Kripke's puzzle. The author presents the case of Pierre, a French speaker, who, on the basis of what he hears about London, says "Londres est jolie", leading us to conclude that he believes that London is pretty. However, one day Pierre moves to London and goes to live in a rather ugly area of the city. He learns English without resorting to translation and is now willing to assent to the sentence "London is not pretty". He is not in a position to equate what he thought of under the name 'Londres' with what he now thinks of under the name 'London' (Kripke 1979, 891-892). Kripke rejects the idea that Pierre has got contradictory beliefs; He says that Pierre lacks information, no logical acumen, and thus he is not able to connect his notion of 'Londres' with his notion of 'London'.

The notion of ‘connection’ used here is one of referential identity: coming to know that the thing *x* referred to as ‘MoP*’ is the same thing *x* referred to by ‘MoP**’.

There is another puzzle, which according to Kripke (1979), arises without substitution. Consider the following utterances:

- (15) Peter believes that Paderewski had musical talent.
- (16) Peter disbelieves that Paderewski had musical talent.

Peter uses the name ‘Paderewski’ for what he takes to be two different individuals. The problem is to explain, given that Peter does not realize that Paderewski the statesman is Paderewski the pianist, how both of (15) and (16) can be true. It may appear that Peter has contradictory beliefs, but this is not the case. Peter is not illogical, he is merely ignorant.

Crimmins & Perry (1989) resolve problems of this kind by noting that the same referent can be associated with two different notions and that failure to connect these notions leads a person to be in two distinct belief states. The most important idea of this chapter is that a belief report is about an **unarticulated constituent**, a propositional constituent that is not linguistically articulated. Context is what leads to the specification of this constituent:

We shall say in such cases that the notions that the belief report is about are *provided* by the utterance and its context (Crimmins & Perry 1989, 975).

5. Bach’s view of belief reports.

Let us now return to the issue of belief reports, armed with the notion that pragmatic intrusion furnishes a fully truth-evaluable proposition. Given that it is implausible that the mode of presentation is furnished by the semantics of the belief report, we conclude that it must be supplied by recourse to pragmatics. Now we must proceed in the direction of a pragmatic theory of belief reports.

Bach’s view of belief reports is one such theory. Bach says:

(...) ‘that’-clauses are not content clauses. The specification assumption is false: even though ‘that’-clauses express propositions, belief reports do not in general *specify* things that people believe (or disbelieve) – they merely

describe or characterize them. A ‘that’-clause is not a specifier (much less a proper name, as is sometimes casually suggested) of the thing believed, but merely a descriptor of it. A belief report can be true even if what the believer believes is more specific than the proposition expressed by the ‘that’-clause used to characterize what he believes (Bach 2000, 121).

This quotation is important because it stresses the role played by **pragmatic intrusion** in fleshing out the proposition believed on the basis of the surface elements appearing in the ‘that’-clause and the context which serves to enrich or expand the proposition.

Let us see how this approach allows us to handle the Paderewski case. Let us recall that the problem arises due to **ignorance**: Peter has two notions of Paderewski, which, as Crimmins & Perry say, he is not able to connect. Peter believes of Paderewski, of whom he knows that he is a musician, that he has musical talent. And he believes of Paderewski, of whom he knows that he is a statesman, that he has no musical talent. Peter is not able to connect the two notions as he is ignorant of the fact that Paderewski the statesman is nothing but Paderewski the musician. Thus, we can say of Peter both (17) and (18):

(17) Peter believes that Paderewski had musical talent;

(18) Peter disbelieves that Paderewski had musical talent (understood as: Peter believes that Paderewski had no musical talent).

(17) and (18) do not attribute contradictory beliefs to Peter – even if the speaker of (17) and (18) knows that there is a referential identity between the two instances of ‘Paderewski’ – provided that Peter is not able to connect the two notions he has of Paderewski.

Bach explains why (17) and (18) are not contradictory statements by saying that the ‘that’-clauses do not fully specify the propositions believed (by Peter), but simply **characterize** them. To fully specify what Peter believes in the two cases, we need to flesh out the proposition corresponding to the ‘that’-clause, using **appositives** both in (17) and (18), thus obtaining (19) and (20):

(19) Peter believes that Paderewski, the pianist, had musical talent;

- (20) Peter disbelieves that Paderewski, the statesman, had musical talent.

Bach writes:

This difference could be indicated by using the appositives ‘the pianist’ and ‘the statesman’ after the name ‘Paderewski’. Using one appositive rather than the other would be sufficient in the context to differentiate one belief from the other, although both beliefs are such as to be true only if Paderewski had musical talent (Bach 2000, 126).

Bach’s proposal is a step forward towards a more accurate theory of belief reports. He proposes to make modes of presentation of the referent ‘Paderewski’ explicit, in line with Salmon’s idea that:

The important thing is that, by definition, they [modes of presentation] are such that if a fully rational believer adopts conflicting attitudes (such as belief and disbelief, or belief and suspension of judgement) toward proposition p and q, then the believer must take p and q in different ways, by means of different modes of acquaintance, in harbouring the conflicting attitudes towards them – even if p and q are in fact the same proposition (Salmon 1990, 230).

Bach’s idea that an appositive qua-clause is supplied by pragmatics in utterances such as (17) and (18) is shared by Bezuidenhout (2000). Like Bach (2000), she too claims that this process of inferential enrichment is the norm, rather than an ‘ad hoc’ way of resolving a puzzle. Bach’s idea that modes of presentation must be made explicit in order to deal with Paderewski-like cases is important. Yet, something else has to be said.

First, (19), (20) are interpretatively ambiguous: on one interpretation the appositive specifies how Peter thinks of Paderewski; on the other, it specifies who the speaker of the entire sentence has in mind². The interpretative ambiguity must be resolved by pragmatics, as I argue in a later section, responding to a similar objection to my analysis of belief reports based on pragmatic intrusion and

² This observation was made by a commentator.

specifically on the provision (or specification) of null appositives. The problem is tackled later on in this chapter.

Second, the fact that Peter has two different notions of Paderewski, as is known in context, does not suffice to make his two beliefs about the two notions of Paderewski's non-contradictory. Suppose, in fact, that at some stage he thought that the two notions of Paderewski could be connected to the same individual (and this realization can be grasped pragmatically by using further appositives). Then, at least for a moment, he must have had contradictory thoughts contemplating the thoughts (19) and (20). So, I think that the 'that' clauses must be specified further and that the inexplicit hidden constituents to be fleshed out must include Paderewski, the pianist, and individual distinct from Paderewski, the statesman.

The most interesting part of the theory is that it is applicable to **all** substitution cases, as well. Thus, in principle it ought to explain why is it that a speaker can say both of:

- (21) Alexander believes that Cicero was a great orator of the past;
- (22) Alexander does not believe that Tullius was a great orator of the past

provided that it is part of contextual knowledge that Alexander does not know that Cicero is also known under the mode of presentation 'Tullius'.

Bach does not go to great length to explain how pragmatic intrusion can account for examples such as (21) and (22). Intuitively, it is clear that what makes the two statements non-contradictory is the fact that Alexander is not able to connect his two notions 'Cicero' and 'Tullius'. Now, pairs of sentences such as (21) and (22) must always be evaluated in context, where contextual knowledge provides linguistic materials that expand the that-clauses further and make it clear that the notions 'Cicero' and 'Tullius' are not linked. The missing constituent is something like (23) or (24):

- (23) the great orator;
- (24) the man I bumped into yesterday at the market place

6. On modes of presentation again! (pragmatic intrusion).

I think it is time to explain opacity phenomena and the puzzles associated with belief reports through a theory of pragmatic intrusion in which modes of presentation of propositions are supplied through pragmatics.

Suppose for a minute that a pragmatician says that pragmatics is merely involved in explaining why we find it misleading to report a belief utterance by substituting an NP with a coreferential one, while he accepts that the beliefs reported are essentially the same (thus excluding pragmatic intrusion into the proposition expressed). Now, consider a sentence such as (25):

(25) John believes that Mary Smith is clever.

Suppose that the referent of 'Mary Smith' is x ; then it would be reasonable to assume that if (25) is true, John must believe of x , under the mode of presentation 'Mary Smith', that she is clever. For Salmon (26) expresses the same proposition as (27):

(26) John believes that she is clever;

(27) John believes that Mary Smith is clever.

He explains the fact that an ordinary speaker surely finds that (26) and (27) have distinct truth-values by resorting to a Gricean pragmatic reasoning. If a speaker attributes the pronominal mode of presentation 'she' to John, it would be misleading to use a more informative sentence such as (27), which leads the hearer to attribute the mode of presentation 'Mary Smith' to John. Thus, it is not really reasonable to trust the ordinary speaker's judgements, who cannot distinguish between truth-conditional and non-truth-conditional elements of meaning.

My view instead is that Salmon's pragmatic view must be refined and recast in terms of Relevance Theory in order to reply to some obvious objections.

Consider a simple sentence such as (28)

(28) John believes that Mary Smith is clever.

(28) has the following logical form:

John believes of x that she is clever.

Pragmatics adds the constituent: under mode of presentation Mo/Mary Smith. Thus, via pragmatic intrusion, we have:

John believes of x , under MoP, that she is clever.

A more elegant representation of this interpretation is certainly the following, adapted from Green (1998):

BEL [John, that Mary Smith is clever, ft (John, 'Mary Smith is clever')]

Where $ft(x, S)$ is a function that takes a person x , a sentence S and a time argument t as arguments and gives as values the way x would take the information content of sentence S , at t , were it presented to him or her through the very sentence S .

I now try to provide an explanation of the interpretation of belief reports on the basis of Sperber & Wilson's relevance theory (Sperber & Wilson 1986). The principle at work in the pragmatic specification of modes of presentation is the following:

Communicative principle of relevance. According to this principle, every act of ostensive communication communicates a presumption of its optimal relevance. An ostensive stimulus is optimally relevant iff it is (a) relevant enough to be worth the hearer's attention; (b) the most relevant stimulus the speaker could have produced given her abilities and preferences (Sperber & Wilson 1995, 270).

The sentence (28) is optimally relevant if the NP 'Mary Smith' is the mode of presentation associated (pragmatically) with the referent of 'Mary Smith', in other words if it plays some role in the identification of reference for the believer. Given that relevance is a ratio of contextual effects and cognitive efforts, it goes without saying that the use of a proper name in the that-clause of a belief-sentence is maximally relevant if it has maximal positive effects, in other words if it does not

just provide a referent but if it is actually used by the believer in identifying the referent in question.

Now we can explain in a more articulated manner why we have the intuition that (26) and (27) do not have the same truth-conditions, given that we accept that pragmatic intrusion contributes to a fully truth-evaluable proposition (e.g. Carston 1999) and, thus is part of what is **said** (according to Carston's 2002 notion of what is said, not according to Bach's notion of what is said; on the distinction see Burton-Roberts 2005 and Burton-Roberts 2006). Furthermore, in Capone (2006a) and Capone (2009) I also argued that pragmatic inferences that contribute to pragmatic intrusion are not cancellable (see also Capone 2009; Burton-Roberts 2005; 2006 in support of this view). If my ideas are correct, the fact that (26) and (27) intuitively have got distinct truth-conditions (ordinary speakers would perceive them to have distinct truth-conditions, regardless of how things are from a theoretical point of view) is just the consequence of our theoretical assumptions about explicatures: explicatures are non-cancellable. My ideas are in line with Sperber & Wilson's view of explicitness outlined in *Relevance* (1986, 182):

Explicitness:

An assumption communicated by an utterance U is explicit if and only if it is a development of a logical form encoded by U.

On the analogy of 'implicature', Sperber & Wilson call an explicitly communicated assumption **an explicature**. Logical forms are 'developed' into explicatures by inferential enrichment. Every explicature, then, is recovered by a combination of decoding and inference.

The picture we have come to is somewhat different from that adopted in other pragmatic views of propositional attitudes. Salmon (1986) would say that the sentences (26) and (27) have got the same truth-conditions, because he essentially leaves pragmatic intrusion out of the picture. Or, to be more precise, he allows pragmatic intrusion up to a point, until the referents of 'she' and 'Mary Smith' are made part of the interpreted logical form, but does not accept (at least not explicitly) a more radically intrusionistic view, like the one I proposed along the lines of Carston (1999) or Wilson & Sperber (2002), in which the provision of modes of presentation is made part of the proposition uttered. Consider (29) and (30):

- (29) John believes that Hesperus is Hesperus;
 (30) John believes that Hesperus is Phosphorus.

Ordinary speakers appear to attribute distinct truth-conditions to these statements. Salmon explains the oddity in such a judgement by saying that it would be misleading for a speaker who commits herself to (29) to utter (30), since the reference to a mode of presentation (of the reference) is part of the pragmatics of the belief report. Yet, he does not consider the case in which (29) and (30) are not real utterances, but just thoughts, to be attributed to a thinker (in silent utterances). In this case, his pragmatic strategy is not available. Yet, Sperber & Wilson's theory helps to explain why (29) and (30) are distinct thoughts by providing each of them with a distinct explicature.

Timothy Williamson (2006) makes an interesting comment. After all, even in thought one can entertain a proposition under the guise of a sentence. Presumably, Williamson thinks that modes of presentation, in thought, are supplied directly by the sentence used to express that thought. But this not entirely persuasive. Even in thought, we need a mechanism that blocks substitution in opaque contexts, so that we shall not say of Mary who thinks that John believes that Hesperus is Hesperus that she thinks that John believes that Hesperus is Phosphorus. As Williamson argues, the sentence used in thought surely can provide a suitable mode of presentation of the reference, but it is no guarantee that substitutions of synonymous expressions is blocked. This is blocked, instead, by mechanisms of interpretations such as those advocated by relevant theorists. A sentence can provide a suitable mode of presentation to a thought, but the principle of relevance ensures that that mode of presentation is the only one under which the proposition is held by the believer or, if this is too strong, that that proposition is not necessarily held under all possible modes of presentation of the reference.

It could be added that surely implicature can occur in the realm of thought as well: I think to myself "Good thing I took the medicine and got better!". I am sure that I take the content of my thought to be some such proposition as 'good thing I took the medicine and, as a result, got better', even though I am not communicating with anyone else.

I quite agree with the above. Pragmatic processes are present even in thought, provided that thought occurs through some linguistic sentences. I take the remarks above as supporting my view that pragmatics is needed to construct propositions even in thought and, thus, to provide modes of presentation. Presumably, what I was opposing is a view that accepts that it would be misleading for a speaker to utter a belief sentence in case the NP in the embedded that-clause is not a mode of presentation under which a referent is thought of by the believer (see also Recanati 1983). Such a view is intrinsically connected with the notion of what goes on in conversation, while we need an account that dispenses with the representation of a real conversation, because belief sentences are employed even in thoughts. Of course it is right that even in thought there are inferences which are the counterparts of the implicatures which would be triggered in actual communication, but here I assume a relevance theory explanation of sentences-in-thought such as “Good thing I took the medicine and got better!”, as this is more in tune with the issue of sentences-in-thought. Gricean explanations are more suitable for actual communication acts.³ In fact, the way Levinson’s (2000) Q- and I-Principle are formulated seems to me to need some notion of actual communication (The Q-principle says that we should not proffer an assertion that is weaker than our knowledge of the world allows, unless asserting a stronger assertion violates the I-Principle. The I-Principle says that we must produce the minimal semantic clues indispensable for achieving our communicative goals (bearing the Q-principle in mind)).

Furthermore, it can be evinced from Green’s (1998) discussion of Salmon’s treatment of belief reports (perhaps one of the clearest expositions of Salmon’s views) that the Gricean explanation of belief reports makes a heavy use of the notion of ‘asserting’, while my relevance theory explanation of the phenomena in question makes allowance for sentences in thought.

Unlike Salmon, I adopt a fully intrusionistic picture in lines with Carston (1999) and Wilson & Sperber (2002), and say that the propositions which John is said to believe in (29) and in (30) are distinct, as they include distinct modes of presentation. The pragmatic machinery is responsible for the fleshing out of the propositions believed and the inclusion there of distinct modes of presentation.

³ To deal with this issue exhaustively one needs one further paper.

It might be worth our while examining more closely the way Relevance Theory can deal with belief reports, in the light of a natural objection arising from accepting Devitt (1996).

Sperber & Wilson (2004) propose the following sub-tasks in the overall comprehension process

- (a) Constructing an appropriate hypothesis about explicatures by developing the linguistically-encoded logical form;
- (b) Constructing an appropriate hypothesis about the intended contextual assumptions (implicated premises);
- (c) Constructing an appropriate hypothesis about the intended contextual implications (implicated conclusions).

So, let us reconstruct the steps required in processing belief reports of the type:

- (31) John believes that Mary is pretty.

The speaker uttered (31) saying that John believes the proposition that X is pretty. The speaker could have chosen a range MoP1, MoP2, MoPn of modes of presentation to present John's belief (about Mary), but he chose 'Mary' as a mode of presentation of X. The utterance (31) comes with a presumption of optimal relevance, that is with the promise that the actual linguistic choice is determined by the intention of causing maximal contextual effects with minimal processing costs. The hearer now realizes that the reason why 'Mary' was chosen in the utterance (31) is that the speaker thus hopes to obtain maximal relevance by increasing contextual effects. The interpretation according to which 'Mary' is the mode of presentation under which the belief is held involves maximal positive effects because it serves to differentiate what John believes from what John does not believe, or at least it serves to specify a more fine-grained ascription of belief⁴.

A natural objection is that an utterance of (32)

- (32) John believes that Mary Smith is pretty,

⁴ I was told that the choice of NP may ease the comprehension process (thus reducing processing costs) and that reduced effort may increase overall relevance.

can be interpreted without having to assume that ‘Mary Smith’ is a mode of presentation under which the belief is held. After all, as Devitt (1996) says, the NP could be used to facilitate recognition of the referent to the hearer of the belief report, in which case it need not play a crucial role in the mental life of the believer (John). Presumably, Devitt’s position is in line with Quine (1960, 218):

Commonly the degree of allowable deviation depends on why we are quoting. It is a question of what traits of the quoted speaker’s remarks we want to make something of; those are the traits that must be kept straight if our indirect quotation is to count as true. Similar remarks apply to sentences of belief and other propositional attitudes (Quine 1960, 218).

Now, I do not deny that there might be a context in which the hearer H, faced with (32), replies: “Sorry, I do not know Mary Smith”, and then the speaker replaces (32) with (33):

(33) John believes that [our department’s secretary]₀ is pretty.

Well, in this case, given that the context is different, and that the hearer understands that the correction has been made to enhance the hearer’s comprehension, maximal relevance is achieved if ‘our department’s secretary’ is not the mode of presentation under which the belief is presented (to the believer). In a context in which the focus is on action, maximal positive effects are achieved if one uses descriptions which facilitate the action in question. The **practical** concerns which lie at the heart of Devitt’s treatment do not necessarily clash with my view, since Devitt must be aware that his proposal is based on heavy contextual assumptions. What, nevertheless, I would like to stress is that Devitt’s treatment does not do justice to the standard pragmatic interpretation of belief reports. After all, the use in (33) is not perceived to be the normal, ordinary use of belief reports, which is to throw light on the mental life of the believers. There is one more thing to be added. In the sentence (33), I think there is an implicit mode of presentation which I marked as 0, which is bound (through pragmatic anaphora) to the NP ‘Mary Smith’ in (32). This is not to suggest that there is always this implicit mode of presentation in the structure of the

explicated thought, yet an array of implicit contextual assumptions may make the interpretation of this 0 as a neutral (and inert) mode of presentation under which the belief is held by the believer.

To consider an example adapted from Devitt (1996), suppose that my cousin, Robert McKay, has recently murdered John Gruff. I know that he is the murderer (furthermore suppose that he always tells me what he does). We happen to read the local newspaper, which has published an interview with an important detective, Sherlock Holmes. The detective provides some details about the state of the investigation and says that he is far from knowing the identity of the murderer. Among the things Sherlock Holmes says is that he believes that the murderer is insane. So both my cousin and I know that Sherlock Holmes is far from knowing the name of the murderer. Yet, I say

(34) Sherlock Holmes believes that the murderer is insane. Thus, Sherlock Holmes believes that [Robert McKay] 0 is insane

where ‘Robert McKay’ is the mode of presentation adopted to make Sherlock Holmes relevant to H and to induce him to reflect on his mental state and 0 is the mode of presentation under which the belief is originally held by the believer.

That contextual assumptions must be taken into account in pragmatic interpretations is well-known. I do not take these as fatal objections to my relevance-theoretic treatment of belief reports. Green (1989), instead, believes that cases such as the one by Devitt militate against a pragmatic analysis of belief reports, presumably because he would like to align inferences such as the ones arising from belief reports to almost-universal implicatures such as those arising from utterances of “I lost a contact lens”. Green argues that the implicature “I lost my contact lens” falls under the scope of negation and of modal embedding (conditionals) and, thus, is an ideal candidate for inclusion in “what is said” by a speaker. Implicatures from belief reports lack the almost-universal feature, presumably because they are defeated in some contexts. Yet, Green undervalues cases of defeasibility such as “I found a contact lens” where the intuitive understanding is that the speaker found somebody else’s contact lens. Thus it can be doubted that there are near-universal implicatures in Green’s sense. It appears to me that what makes inferences of belief reports eligible to be part of what is said, in addition to being part of the proposition

expressed, is the fact that they are not cancellable. I personally find the examples of cancellability by Green (1998) based on Barwise & Perry (1981) and Berg (1988) hard to swallow. Furthermore, the fact that, as Green notes, ordinary speakers' judgements consider utterances of belief reports that express the same proposition but contain distinct modes of presentation as truth-conditionally distinct seems to militate against the view that the inferences of belief reports are implicatures and in favour of the view that they are part of what is said and of (uncancellable) explicatures.

So far, I have argued that pragmatic intrusion is responsible for enriching the logical forms of belief sentences and fleshing out the full truth-evaluable propositions associated with belief reports. The processes I have inquired into are **unreflective**, and largely intuitive, in line with considerations by Wilson (2000, 417). The fact that explicatures are mainly unreflective can explain why it is that most of us are inclined to think that belief reports are not interchangeable *salva veritate* if an NP is replaced with a coreferential expression.

Devitt's (1996) approach may be seen as an ideal candidate for the treatment of belief reports presumably because it has the merit of reconciling Millian with Fregean theories (Davis 2005). Yet, the approach is unsatisfactory because it does not address semantic and syntactic problems properly. The way Devitt hopes to reconcile both positions is to say that each NP (or AP) within the clause embedded in a belief verb expresses both a referent and a mode of presentation. Yet this apparently conciliatory move does not take into account the syntactic difficulties which were a threat to Schiffer's theory. Surely Devitt would not want to say that each NP (or AP) **semantically** expresses both a referent and a mode of presentation. Such a claim, even if possibly true, does not explain the opacity problem: the fact that belief contexts block the application of Leibniz's law. Consider Leibniz' Law:

Two things are identical with each other if they are substitutable preserving the truth of the sentence (Jaszczolt 2005, 120)⁵.

⁵ Williamson (2006) correctly argues that a better formulation of Leibniz's law is required. The reader is referred to Asher (2000) and in particular to his identity principle: Suppose that ϕ is an expression denoting an abstract entity, that ϕ contains an occurrence of a name α , and that the denotation of α is the same as the

Arguing in favour of the (semantic) association of every NP with a referent and a mode of presentation would *ipso facto* create a problem in that Leibniz's law would then be inapplicable even in the case of NPs outside the scope of belief-like operators (opacity would be exported outside the scope of belief verbs); not to mention the fact that not all NPs can be directly associated with referents (what about 'beauty', 'wealth', 'justice'?).

An additional problem is that, in the spirit of his conciliatory proposal, Devitt grants that both transparent and opaque interpretations are licensed by belief reports, following Quine (1960). He grants that a sentence such as (35)

(35) Ralph believes that Orcutt is a spy

has got the two following interpretations:

(36) Ralph believes of Orcutt that he is a spy;

(37) Ralph believes that (assents to) "Orcutt is a spy".

(36) constitutes the transparent construal, whereas (37) constitutes the opaque construal.

A thorny problem for Devitt (and for Quine) is that, everything being equal, a **univocal semantic representation** should be preferred to the ambiguity view, on grounds of parsimony (Modified Occam's Razor; see also the important work by Jaszczolt 1999, who tries to eliminate ambiguities in favour of univocal interpretations). Another problem would be that the transparent reading, given his general conciliatory strategy of associating an NP both with a referent and with a mode of presentation, should be obtained by suppressing the mode of presentation the referent is associated with. So, Devitt faces the hard task of explaining where the mode of presentation comes from (semantically) in the opaque construal; in addition, he must explain how the mode of presentation is suppressed in the transparent construal. The move of resorting to the context of utterance is not

denotation of β , then the denotation of $(\varphi) = \text{denotation}(\varphi [\beta/\alpha])$. Something along these lines is required.

allowed him, if his strategy is not pragmatic, but merely semantic. His strategy is clearly not pragmatic, given what he says, because he invokes no pragmatic machinery to explain what he assumes. In my opinion, explaining in semantic terms how a mode of presentation within the scope of 'believe' is associated with an NP involves syntactically deriving the mode of presentation from the belief verb and claiming that it is an argument of the verb. But this move is not devoid of problems, as Schiffer and Recanati convincingly noticed.

A pragmatic approach avoids the proliferation of senses (the ambiguity problem) and also explains why in some contexts, but not in others, modes of presentation are suppressed. It also explains why the opacity construals are *default*⁶, achieved by maximizing relevance. The transparent interpretation is simply achieved by preventing a mode of presentation from arising, and thus needs a context in which the suppression of the mode of presentation is mandated by background knowledge. I also want to consider that the suppression of a mode of presentation should not be considered a case of defeasibility. Explicatures are derived/constructed through unreflective pragmatic mechanisms that promote the most relevant interpretation, that is to say, the one which has the greatest amount of positive cognitive effects. An interpretation which reduces the possibility of mistaken action is more relevant than an alternative interpretation (in a context in which action is being focused on) because it maximises contextual effects. This is why MoPs are prevented from arising in certain cases in which there is a heavy emphasis on the facilitation of action (presumably in Devitt's cases discussed above).

My proposal has much to do with modes of presentation, but in a sense, it ignores a very important fact pointed out by Jaszczolt (1999). An NP embedded in that-clauses in belief reports has the main function of referring to an entity that belongs to the real world. In my approach, this important fact can be reconciled with the fact that the referent is normally associated with a specific mode of presentation. That NPs within that-clauses of belief reports refer to **extra-linguistic entities** is also ensured by inferential pragmatics, since a that-clause which provides information about the world, in addition to providing information about the believer's mental life, is more informative, as it **eliminates** a greater amount of states of the world, and is conducive to successful action (on the notion of informativeness, see the important

⁶ I am not arguing that the inferences in question are the result of default rules, but only that they standardly get through.

work by Levinson (2000), as well as his papers on anaphora). Thus relevance is increased, as a result of increased contextual effects.

7. Further considerations on null appositives

What I have so far proposed is that pragmatic intrusion provides a specific mode of presentation, while I have suggested that there is a constituent present in the structure of the explicature, which has the features of a pronominal or a free variable and is a (pragmatic) empty category, in that it does not receive a phonological representation. This is a **null appositive** (it should be clear, however, that I am not proposing that the variable is present at LF). The possibility of an NP's having an appositive is exemplified by sentences such as (38)

(38) Mary, the President of our union, is clever.

An appositive is surely a modifier, in that it adds further qualifications or restrictions to those expressed by the main NP. The appositive adds a superior node to an NP node, a node which has similar features, thus an NP. Of course, it is important to know whether the appositive adds a further constituent to the main proposition expressed by the sentence. My answer is that it does not and simply makes the referring potential of the name it is an appositive to more explicit. Given that, in our belief sentences, the appositive representing a mode of presentation is a null element, we can represent it in this way:

(39) John believes that [NP [NP Mary] [NP 0]] is clever.

The most thorny problem I can see, with this proposal, is that, after all, we would have to generalize it to all NPs. Thus, in (39), the NP 'John' too would have to be associated with a null Mode of Presentation. Is not then opacity created in subject position too? Well, my proposal crucially hinges on the interaction between the empty category 0 and the verb 'believe' which has 0 in its scope. 'John' is outside the scope of 'believe' and thus no interesting interaction, resulting in an opacity effect, obtains. The implicit mode of presentation in the subject position of (39) does not result in opacity effects, because it does not prevent substitution *salva veritate* of the NP 'John'.

Another problem with my proposal could be the following. It may be plausibly argued that a sentence such as (38) is truth-conditionally equivalent to (40):

(40) Mary, who is the President of our union, is clever.

Now, suppose we embed this sentence into a belief sentence, we obtain:

(41) John thinks that Mary, who is the President of our union, is clever.

However, on one interpretation the relative clause just gives more information about Mary without shedding light on how John thinks of her. The objection is a natural one and a very good one too. I pointed out in a previous section that Bach's use of appositive clauses could not be the end of the story and that further pragmatic processing was involved. The objection is one further reason for assuming that a story in terms of Relevance Theory is needed. This must be a story that explains why an appositive clause is needed and how this appositive clause is assumed to be part of the believer's mode of presentation of the proposition believed. The problem, however, one might retort, arises not really from implicit appositives, but from explicit ones, like the ones above, which can be understood '*de re*' (as not being part of the believer's mode of presentation). Cases such as this can be disposed of with a pragmatic treatment in line with Sperber & Wilson's considerations. The implicit appositive clause is interpreted as providing a mode of presentation that serves to specify further the believer's belief state because this interpretation has greater contextual effects than a merely referential interpretation. In other words, considerations analogous to the ones I adopted in resolving Jaszczolt's problem apply here. This problem seems to me of great theoretical importance. I will take it up later.

The discussion so far has hinged on the assumption that we can have something like null appositives, specifically modes of presentation, in the structure of NPs belonging to that-clauses embedded in verbs of belief. The literature on pragmatic intrusion is characterised by endless discussions on whether we should posit empty constituents in logical forms of sentences such as "It rains". Recanati (2004) is a champion of the view that we should not posit these empty categories at logical form. I must say, in the words of Mey (personal communication) that in cases such as the

one discussed by Recanati there is not clear-cut evidence in favour of one or another theory. But have we got independent evidence for the existence of this null appositive that modifies NPs within the scope of belief verbs? I propose that we should set aside the task of assigning null appositives at logical form and remain content with stipulating that such appositives appear in the propositions expressed (so we are following Recanati 2004 and Carston 2002). Some crucial and indubitable evidence comes from sentences such as (42):

(42) John believes that Mary Pope went to Paris and that she had fun.

I propose we should analyse (42) as providing evidence for a proposition such as:

(43) John believes that [Mary Pope]₀ went to Paris and that [she]₀ had fun.

It is interesting to note that if we allow implicit modes of presentation, we have got a double anaphoric pattern, as the indexes show:

(44) John believes that [Mary Pope]_i _{0n} went to Paris and that [she]_i _{0n} had fun.

The subscript *i* represents the reference of ‘Mary Pope’, and this is attributed through coindexation to ‘she’. Instead, *n* is the subscript attributed to ₀, the mode of presentation associated with ‘Mary Pope’ (which must be coindexed with the form ‘Mary Pope’), and that is coindexed with the implicit mode of presentation ₀ associated with ‘she’. Notice that, unless we have got this (conceptual) anaphoric chain, which is possible only through the existence of null modes of presentation (or null appositives), it would be possible to understand (42) allowing ‘she’ to be intersubstitutable with any NP at all that has the same referent as ‘Mary Pope’, with no regard for the mode of presentation ‘Mary Pope’. But this is not the natural interpretation of the utterance.

Further evidence comes from control structures:

(45) John believes Mary Pope to be in Paris and [PRO]₀ to be working hard (instead she is having fun with her other boyfriend).

The control structure ensures that the reference of ‘Mary Pope’ is transmitted through anaphora to PRO; however unless we posit that PRO has got a null appositive in the explicature, we cannot account for the opacity of the structure, as certainly by replacing PRO with an NP coreferential with ‘Mary Pope’ (but distinct from it) a statement having different truth-conditions is obtained⁷.

A more interesting piece of evidence comes from Italian control structures:

(46) Maria crede di PRO essere intelligente.

(lit. Maria believes PRO to be intelligent)

Maria believes she is intelligent.

Suppose we call Maria ‘Maria’, but she does not know that this is her name; in fact, she does not know that she has a name. Maria thinks of herself under some mode of presentation of the self (a first-person mode of presentation), but this does not include the name ‘Maria’, which, instead, is the mode of presentation we associate with her. This case strongly supports the idea that we must posit a propositional structure such as the following:

(47) Maria crede di [PRO] 0 essere intelligente.

In fact, while PRO in the present case receives its reference through an anaphoric link with Maria, it cannot be associated with the mode of presentation ‘Maria’. We thus need a way of signalling that PRO must be possibly distinct from 0 and that 0 must be possibly distinct from ‘Maria’. 0 is associated with PRO, but not through anaphora, only as a null appositive, which is capable of having the meaning of ‘whatever coincides with the subject of belief’⁸. This is what Lewis (1979) calls an attitude ‘de se’. Higginbotham’s (2004) considerations on the “internal” aspect of PRO (in the context of a discussion of ‘de se’ beliefs) are applicable here: what is believed by Maria to be intelligent is the subject of the experience BELIEVING (see also Stanley & Williamson 2001, for an analogous view). The example (47) is

⁷ In this respect, my view is different from Salmon’s.

⁸ My example is reminiscent of an example by Stanley & Williamson (2001), who actually use a case of amnesia to exemplify *de se* interpretations.

reminiscent of Castañeda's (1966) famous example "The editor of Soul knows that he is a millionaire". This, according to Castañeda, does not entail that the Editor of Soul knows that the Editor of Soul is a millionaire.

In Capone (2000), I proposed that clitics in **clitic-doubling constructions** qualify propositions believed or known in a special way, anchoring them to discourse. While I recognize that more than one view is tenable in connection with the Italian clitic 'lo', I now think that a theory along the lines I have just proposed is not untenable. Consider (48):

- (48) Giovanni lo sa che Maria è a Parigi.
 (lit. Giovanni it knows that Mary is in Paris)
 (Giovanni knows that Mary is in Paris).

I claimed that the clitic 'lo' qualifies the proposition known as being part of shared knowledge. I have wavered between a semantic and a pragmatic analysis, but this is not important for the present chapter, as the semantic analysis may be seen as deriving diachronically from the pragmatic analysis. In any case, both Uriagereka (1995) and I have claimed that the clitic has the effect of qualifying the that-clause embedded in the verb of propositional attitude as being part of the commonground. Yet Uriagereka claims that the that-clause is an appositive to the clitic, the clitic functioning as an argument to the verb. This view is not implausible. Suppose, however, that we reverse this view and claim that the clitic is an **appositive** to the that-clause. In this way we obtain an analysis parallel to the one I have proposed for verbs of propositional attitude in general. Modes of presentations are appositives. In the case of clitics (in clitic-doubling constructions), they are pronominal appositives, which qualify a proposition in a certain way: they can function as words that imply certain modes of presentation. Of course, they need not be called modes of presentation, except in a special sense. The word 'lo' qualifies the embedded proposition (in clitic doubling constructions) as being part of the commonground, but the word itself is not the mode of presentation of the embedded proposition. It implies a mode of presentation.

Further evidence comes from verbs of propositional attitude like 'want', which mandate control structures. Consider (49):

(49) Mary wants Cicero PRO, but not Tullius, to come to her party.

Suppose Mary is not aware that Cicero is Tullius; she would like Cicero to come to the party, but she would like Tullius not to come to her party (say she has always heard nice things about Cicero but bad things about Tullius). If she knew that Cicero is Tullius, she would not let him come. (49) ought to be analysed as (50):

(50) Mary wants Cicero [ti 0 to come to the party] but she does not want Tullius [tii 0 to come to the party]⁹.

Unless we posit null appositive modes of presentation associated with ti and tii, the sentence (50) has to be perceived as a contradiction, since ti inherits its reference from ‘Cicero’, tii inherits its reference from ‘Tullius’ (in conjunction with the premise that Cicero = Tullius). However, the sentence is not contradictory because Mary, due to her frame of mind, will not admit Tullius to her party (given the bad things she has heard about him), but she will allow Cicero (given the nice things she has heard about him). The sentence (50) reminds us of Carston’s famous example (originally from Cohen 1971):

(50) If the king of France dies and France becomes a republic, I shall be happy, but if France becomes a republic and the king of France dies I shall be unhappy.

The sentence (50) has an appearance of contradiction, unless we furnish the explicatures and provide a temporal and causal interpretation for the material in each if-clause. When the underdeterminacy is resolved, we no longer face a possibly contradictory sentence (however, see Capone 2006 for a more detailed proposal, in addition to Burton-Roberts 2005). Analogously, by furnishing the explicature of (50) and in particular by coindexing each 0 (implicit mode of presentation) with the NP 0 is an appositive to, we obtain a statement which is no longer contradictory.

My proposal abides by Carston’s precept that:

(...) pragmatic processes can supply constituents to what is said solely on communicative grounds, without any linguistic pointer (Carston 2002, 23).

⁹ I am adopting this analysis from Carnie (2002).

There is an interesting interaction between modality and belief reports, and it might be useful to analyse the modal ‘ought’. Obligations are normally imposed by the rules of society on people who act in certain roles. Were they not to act in such roles, they would not contract the relevant obligations. For example, as I am a teacher I ought to lecture almost every morning in my school (but not on Sundays!). The mode of presentation is important, as the obligation is perceived only under a specific mode of presentation, which includes a public role (it follows that I should not lecture on Sundays, when I do not act as a public official). So consider the following example:

(51) John believes he ought to lecture in the morning.

Although some may deny this, I believe that some implicit modes of presentation are at play here, explaining that the propositional form includes that John believes that he, qua teacher, ought to lecture in the morning (every working day of his life as a teacher, but certainly not when he retires); in other words, the modal ‘ought’ selects a **stage** of John’s life in which he is a teacher and cannot be extended to periods of time in which he is retired). If such considerations make sense, we can explain why substitution can be blocked in the following case:

(52) John believes that the Prime Minister ought to go to the ceremony of the opening of the judiciary year;

(53) John believes that Berlusconi ought to go to the ceremony of the opening of the judiciary year.

The expression ‘the prime minister’ in (52) may receive an attributive or a referential interpretation. The attributive interpretation is ‘whoever is the Prime Minister’; the referential interpretation is ‘Berlusconi’. Suppose we confine ourselves to the attributive interpretation. From this interpretation, it follows that if the Prime Minister ought to go to the ceremony, then Berlusconi has to go. Yet, from (52) it does not follow that (53) is unconditionally true. First, the speaker need not know that Berlusconi happens to be the Prime Minister. Secondly, (53) follows from (52) **provided that** the speaker knows of the identity in question and on the

understanding that ‘ought’ is confined to the period in which Berlusconi is in charge as Prime Minister. I think these moral matters ought to have discouraged Devitt (1996), who, instead, uses the attributive/referential interpretative ambiguity to argue that substitution of NPs in belief sentences is licit for communicative purposes.

A further example that can be used to prove that implicit modes of presentation do some work at the propositional level is drawn from Seymour’s (1992) paper. Seymour defends a sentential theory of propositional attitudes and essentially believes that a person X believes that S in case he is in a relation R to a certain sentence. If we mention the sentence in order to specify its character, the whole belief sentence reports a belief relation between the agent and a certain linguistic meaning under a mode of presentation that is a certain verbal form. If we mention the subordinate clause in order to specify its content, the agent is then described as being in a relation with the content of the sentence mentioned and, in this case, the sentence no longer behaves as a mode of presentation. Seymour calls the first type of belief **intentional belief** and the second **material belief**. Seymour considers that intentional beliefs (opaque readings) are distinguished by the fact that they are **reflexive**. Material beliefs are distinguished by the fact that they are not reflexive. Thus, if (54) is true, (55) must be true, provided that we consider the opaque reading of (54):

(54) John believes that Mary Simpson went to Paris.

(55) John believes he believes that Mary Simpson went to Paris.

If John is not ready to assent to ‘Mary Simpson went to Paris’, there is no way to derive the inference (55). Seymour’s intuition is most easily explained away by resorting to modes of presentation (given Kripke’s reasonable doubts about the equation of believing and assenting to a proposition). Thus the statement or thought (54) must receive an adequate representation as (56):

(56) John believes that [Mary Simpson]_i 0_i went to Paris

where 0 represents the null appositive or mode of presentation, which is coindexed with the form ‘Mary Simpson’.

Timothy Williamson (2006) says that there is the case in which John lacks self-knowledge or does not grasp the concept of belief. Presumably this amounts to an

objection against my treatment. As a reply, I must consider that in some cases the reflexivity of belief (or of knowledge) comes to the philosopher's help in resolving otherwise insuperable problems. Stanley & Williamson's (2001) discussion of opacity in favour of the idea that knowledge-how is a sub-species of knowledge-that is one such case. Of course, if knowledge-that involved opacity while knowledge-how involved a transparent relation to an embedded proposition *p*, there would be serious trouble for Stanley & Williamson who have taken great pains to analyse knowledge-how in terms of a relation between a (cognitive) agent and a proposition (invoking the semantic machinery of embedded questions). In particular, there would be trouble if there was no significant truth-conditional difference between "Hannah knows how to locate Hesperus" and "Hannah knows how to locate Phosphorus". Stanley & Williamson claim that the latter proposition does not seem to follow from the former. But I think that their analysis tacitly assumes that knowledge is a reflexive relation. Only reflexivity can block the intersubstitution of the two names, because despite the fact that Hannah is able to locate on a map (of the universe) that planet there in the sky regardless of its name(s) (and thus in a sense she is able to locate both Hesperus and Phosphorus), she would not say of herself that she knows she knows how to locate Phosphorus even if she knows she knows how to locate Hesperus. As Seymour would say, it is the reflexive notion of knowledge-how that blocks substitution.

Independent evidence in favour of the presence of modes of presentation of the referent in that-clauses of belief reports comes from what Stanley (2005) says about ellipsis. Stanley argues that explicatures enter into certain linguistic processes such as anaphora and deixis (this is well known since Chomsky 1972, 33). An example such as (57)

(57) The ham sandwich wants his bill now

proves that the pragmatic determination of the referent of 'The ham sandwich' (the person who ordered the ham sandwich) must be available for anaphoric coindexation. The examples of ellipsis are even more interesting. Consider (58):

(58) Bill served a ham sandwich, and John did too.

(58) cannot be interpreted as conveying that Bill served a person who ordered a ham sandwich, whereas John served a ham sandwich. The explicature of the first conjunct of (58) must be available for the understanding of the elided constituent too. In other words it must be used in providing an explicature that reconstructs the missing (elided) constituent.

Analogous considerations apply to metaphorical meaning, which is carried over in ellipsis, showing that explicatures play a role in this linguistic process:

(59) John is a pig and Bill is too.

Now, let us apply ellipsis to belief reports. Consider:

(60) John believes that Kent Clark is not Superman and Fred does too.

We said above that ellipsis carries over the explicature of the first sentence to the elided constituent. Thus, it is not licit to interpret (60) as the thought that John believes a non-contradictory thought while Fred believes a contradictory one. In other words, it is not licit to replace ‘Kent Clark’ with ‘Superman’ in the elided constituent, as a result of a syntactic constraint due to ellipsis: the explicature of the first sentence must be used in reconstructing the meaning of the elided constituent (we could also use talk of a ‘parallel’ explicature). What is this explicature? I assume that it consists in the attribution of the mode of presentation ‘Kent Clark’ to the referent Kent Clark.

Analogous considerations apply to (61) (to use comparatives, which were first taken account of by Carston in her discussion of the semantics/pragmatics debate):

(61) John believes that Kent Clark is better than Superman and Fred does too.

As Carston (2002) and Levinson (2000) noted, the statement ‘A is better than B’ presupposes that A and B are distinct, otherwise it communicates a patently false thought.

Ellipsis in (61) imposes the constraint that the explicature of the first conjoined sentence should carry over to the elided constituent. In particular, the elided constituent cannot be interpreted as “Fred holds the belief that Kent Clark is better

than himself”. What is it that blocks substitution ‘salva veritate’ in the elided constituent? Presumably it is the fact that ‘Kent Clark is associated with the mode of presentation ‘Kent Clark’ in the explicature of the first conjunct and it cannot be associated with the mode of presentation ‘Superman’ in the explicature of the second conjunct. This is best explained away on the assumption that the explicature of the first conjunct is that Kent Clark is associated with the MoP ‘Kent Clark’ in John’s belief and that there is a linguistic constraint due to ellipsis such that the explicature of the first conjunct is carried over to the elided constituent of the second conjunct.

8. An alternative analysis

So far, I have developed an analysis following Bach’s idea that one should posit null appositives (not in the logical form, but) at the level of the propositional form. I have also elaborated on the reasons why null appositives are required at the level of the logical form. Yet, as it could be correctly pointed out, I am not exempted from providing a semantic analysis of the propositional form thus obtained through pragmatic expansion. Furthermore, it is possible that the ideas exposed so far may be further expanded so as accommodate a plausible objection on the part of my hypothetical opponent. Presumably a semantic analysis of the propositional development of null appositive clauses looks like this:

John believes that [NP 0 VP]

where 0 is a null appositive (presumably an NP). This is in line with Del Gobbo’s (2003) idea that appositives expand NPs into NPs (through adjunction). If we adopt the idea that a null appositive is nothing but a relative clause, then we have a further expansion of the structure above:

John believes that [NP [CP who [t is NP]] VP]

where t is the trace of the relative pronoun which moves to a node dominated by CP (complementizer phrase) (see Haegeman 1994).

A plausible consequence of this analysis is that the null relative clauses may be taken as providing the speaker’s mode of presentation of the reference (in addition to the believer’s mode of presentation of the reference). In response to this possible

objection, I think that something else must be said about a sentence such as “John believes that Mary went to Paris”.

Suppose that we exploit the intuitions about the role of clitics in clitic-doubling constructions (e.g. “Giovanni lo sa che p” (Giovanni it knows that p)). We said that the clitic is mainly an appositive (reversing Uriagereka’s analysis according to which the clitic was the complement of the attitudinative and the proposition p an appositive). Furthermore suppose that all attitudinatives have got a null appositive similar in semantic/syntactic structure to the clitic, but syntactically different in that it has got an internal articulation of the type [MoP/SN, MoP/VP], where by MoP we indicate a mode of presentation. We also suppose that this supposition can be held cross-linguistically (unless there is evidence that a language has got a different structure; we are open to the hypothesis that languages may vary according to whether they exhibit tacit or otherwise explicit appositive clauses to attitudinatives).

Well, at this point all attitudinatives have got the following semantic/syntactic structure:

John believes [[that P] [NPMoP VPMoP]]

It is not surprising that there are appositives to sentences, given that sentential arguments are assimilated to NPs. In any case, De Vries (2002) provides a number of examples of appositives to sentences. Given this semantic structure, relevance theory intervenes to supply appropriate binding between NPMoP and the NP occurring in P.

At this point we can easily explain the intuition that non-restrictive relative clauses actually express a speaker’s mode of presentation (in addition to a believer’s mode of presentation). This may find an explanation in the fact that processing effort decreases relevance and, thus, to preserve the high relevance of the non-restrictive relative clause, we need the extra assumption that the non-restrictive relative clause expresses the speaker’s mode of presentation, in addition to the believer’s mode of presentation.

By incorporating qua-clauses not at the level of the uttered sentence, but at the level of the null appositive we have immediately resolved the problem we are tackling.

It may be objected that this alternative analysis places a greater load on the semantics, than on the pragmatics. Yet, we still have to posit appositives to explain the behaviour of clitic-doubling in languages like Italian and to posit null appositives for English is no costly move, given that appositives exist in English and thus the English language must have the semantic and syntactic resources for expressing them.

It may take time to adjudicate between the previous position and the one I express admittedly tentatively in this section. It is not clear to me that the null appositive view, as I expressed it, requires positing free variables (of a complex kind) at Logical Form, in the sense of Stanley (2000), in which case pragmatics is assigned the modest role of filling in these variables, of giving them semantic values by saturating them. The picture I have so far provided is compatible with a full pragmatic intrusion story. The syntactic structure of the null appositive to the embedded clause of a belief report is just the structure of a constituent we mentally supply through pragmatics and it is possibly not part of the semantic structure of a belief sentence.

At this point, we have to clear out how we should treat clitic doubling belief-like constructions. Supposedly, they have a structure similar to the constructions where no clitic appears, with the difference that the clitic is already there, and already functions as an apposition to the *that*-clause. However, there is nothing in the grammar that banishes the idea of having cyclic appositives (or reiterated appositives), an appositive being an appositive to another appositive. At this point, we can assign the following structure to clitic doubling constructions such as “Giovanni lo sa che p”:

Giovanni sa [[che S] [[NP lo] [NP/MoP, VP/MoP]]]

This enables the clitic to provide, by semantic implication, a mode of presentation of the whole of *S*, while [NP/MoP, VP/MoP] supplies a structure which is more articulated and provides modes of presentation variables which can be bound with NPs within the embedded clause. We can think of these variables as present in constituents we assign the sentence through pragmatics. The present analysis assumes, in line with de Vries (2002), that appositives can be stacked in English and in many other languages (an English example is “This man, who came to dinner late, about whom nobody knew anything...”).

9. Loose ends

I think we should now turn to a problem noted by Braun (1998). Braun believes that Fregean theories hold the view that speakers routinely think about other people's mental representations and intend to talk about those mental representations when they utter belief sentences. He thinks that the following case is problematic for Fregean theories. Suppose Gingrich says:

(62) I am a Republican.

Linda hears that and forms a belief about Gingrich. Any of the following is a correct description of Linda's belief:

(63) Linda believes that Gingrich is a Republican;

(64) Linda believes that you are a republican (addressing Gingrich);

(65) Linda believes that he is a Republican (demonstrating Gingrich).

I think Braun's point is to show that, regardless of the mode of presentation Linda makes use of in her (unexpressed) thought, one can use modes of presentation such as 'Gingrich' or pronouns whose content is retrieved relative to the context of utterance. I think Braun concludes that the Fregean emphasis on the mode of presentation of the believer is misguided, as in actual conversation we are in fact more busy in identifying reference.

Wayne Davis (2005, 268) provides the following reply to Braun's argument:

In this case, the fact that 'Gingrich', 'he' and 'you' differ markedly in meaning does not matter at all. But this is clearly a case in which the belief ascriptions are intended transparently. It is on the opaque interpretation that substitutivity fails (Davis 2005, 268).

My answer to Braun is different from Davis's. I think that a point Braun neglects is that in context we know quite well that Linda thinks of Gingrich as Gingrich and thus it follows that a suitable mode of presentation of the referent could either be the name 'Gingrich' (as was done in (63)) or a pronominal, making an **implicit** reference

to the mode of presentation ‘Gingrich’. In other words, Braun’s considerations, valuable though they are, do not militate against my approach which posits implicit appositives (in the explicature) to account for implicit modes of presentation.

Davis (2005) has an example that is useful for the discussion of explicatures (once again I notice that I talk of explicatures, while other theorists talk about implicatures). I believe that explicatures are not cancellable (see Capone forthcoming). Consider in detail Davis’s discussion in terms of the implicatures of belief reports:

The metalinguistic implicature will be cancelled when it is evident in the context that the speaker is using a language not known by the subject being described. If I say, ‘Boris Yeltsin believes that the Pacific Ocean is larger than the Atlantic’, I would not imply or be taken to have implied that Yeltsin believes any English sentence to be true. When we ascribe beliefs to a prelinguistic child or animal, we do not implicate that the subject would believe any sentence at all to be true (Davis 2005, 177).

The hearer knows from the very beginning that Yeltsin’s thoughts are described **through translation** and thus immediately accesses the assumption that the proposition embedded in the belief report is not as fine-grained as if the speaker were talking about an English believer (see also Green 1998 on similar cases). Yet, this pragmatic information amounts to an explicature, not to an implicature. On the present view, the following argument (from Williamson 2006) turns out to be valid (with the proviso that we are talking about coarse-grained beliefs and that, following Kripke, we give up the assumption that believing amounts to assenting to a sentence):

Bush believes that it is raining; Yeltsin believes that it is raining; Bush is not Yeltsin; therefore at least two people believe that it is raining.

The reasoning should suffice to prove that the pragmatically supplied constituents of belief reports (in particular appositive clauses specifying modes of presentation) are part of explicatures. It is propositions that are involved in logical deductive inferences, not fragmentary logical forms. For a full discussion of the role of

pragmatic inference in deductive inferences see Levinson (1983). The explicatures in belief reports like the ones in the Bush/Yeltsin case are motivated by the search for the most relevant proposition effects **compatible with the speaker's abilities and preferences**. The speaker is obviously not capable of translating a belief and of doing justice to the original modes of presentation because the translation would become too laborious and thus relevance would be threatened by undue processing efforts.

I think that the differences between my approach and standard pragmatic approaches (e.g. Salmon's) can surmount an otherwise insuperable objection by Davis (2005, 278), which I deal with below. One who knows the Superman story (I assume everyone is acquainted with it) must accept the following thoughts:

- (66) Lois knows that 'Superman' refers to Superman;
- (67) Lois knows that 'Clerk Kent' refers to Clark Kent.

But in the Millian theory 'Superman' refers to Clark Kent and 'Clark Kent' refers to Superman.

Thus, if we conjoin the thoughts in (66) and (67) with these referential identities, we obtain the following true statements:

- (68) Lois knows that 'Superman' refers to Clark Kent;
- (69) Lois knows that 'Clark Kent' refers to Superman.

But then how can Lois fail to know that the two names are coreferential?

Davis's reasoning is easily surmounted if one accepts, as I proposed earlier, that pragmatic intrusion characterises the proposition expressed. Once the intrusion is fully characterised, the proposition expressed cannot be cancelled. Thus, unless the context is such that Lois knows the identity Superman = Kent Clark, it will not do to replace Superman in (67) with Kent Clark, in that this substitution suggests that the mode of presentation of the reference should be changed too, but this is not possible because the explicature is not cancellable.

Conclusion

In this chapter I have tried to connect two issues which are usually dealt with separately: belief reports and pragmatic intrusion (the semantics/pragmatics debate). I have proposed that advances in the former issue cannot be made without advances in the latter. The semantics of a sentence is often too skeletal to accommodate all the elements of a thought; thus pragmatics must contribute to the expression of implicit constituents. Modes of presentation are usually contributed to a thought through pragmatics. The contribution of this chapter is to explain, through relevance theory, that a belief report opens a window on the mental panorama of a believer and focuses on his way of representing constituents of thought. I also argued that this pragmatic picture is necessitated by linguistic facts, such as anaphora, control, ellipsis, *de se* beliefs, inferential properties of beliefs, modality, presuppositional clitics, etc.

Chapter 7

On the social practice of indirect reports.

Abstract:

In this chapter I postulate a number of principles of language use connected with the practice of indirect reporting, based on the idea that the original speaker should ideally approve of the report in case her words were not completely and arbitrarily transformed. I connect these principles of use to a more general cognitive principle.

Introduction.

This chapter is a contribution to the pragmatics of indirect reports. In this chapter it is my aim to show how semantics and pragmatics can be harmonised in the context of the issue of indirect reports. Semantics involves recursive rules applying to a formal syntax, operates on expressions at a level of 'logical form', and is thus distinguished from pragmatics, through which such logical forms are developed into richer propositions. In many, perhaps almost all, cases a given surface sentence can generate a variety of logical forms. But in a speech context these will be ordered so that some are preferred to others and some are ruled out entirely. Pragmatic principles are mainly responsible for this. If this is right, then semantics and pragmatics are part of a harmonious picture.

In connection with indirect reports, I argue that we need to invoke pragmatic considerations that will complement semantic considerations, semantics and pragmatics working in tandem. Very often, in fact, substituting an NP into an

indirect report results in different speech act being erroneously reported. I do not expect to discover the exceptionless principles governing indirect reporting. Such principles are hard to find. However, what I propose is that such principles are tied to the function of utterances in discourse, and, thus, should be sensitive to the speaker's orientation to the communicative situation.

I propose that the principles governing indirect reports should be ultimately connected with the principle of Relevance by Sperber & Wilson (1986) (the matter whether the principles follow from Relevance Theory's assumptions or simply give independent support to it will remain open for the time being). I therefore choose Relevance Theory as a framework that gives unity to my various considerations.

In this chapter, it is also my wish to draw the implications of the theory of **pragmemes**, which I have advanced in other articles (Capone 2005a, 2004a) following ideas broached by Mey (2001) and to apply them to areas of inquiry which certainly need being revisited with modern and efficient analytical tools, such as the detailed study of the interplay between speech acts and context.

1. Pragmemes and indirect reports.

This chapter owes much to Mey's theory of pragmemes (Mey 2001). For Mey, all utterances are situated and their intended meanings must be recovered starting from the situation of utterance. It should be possible, in theory and in practice, to recover (or discover) the speaker's intended meanings on the basis of the (rich) contextual clues which s/he utilises in getting her message across. This is what Mey says about pragmemes or pragmatic acts:

The theory of pragmatic acts does not explain human language use starting from the words uttered by a single, idealized speaker. Instead, it focuses on the interactional *situation* in which both speakers and hearers realize their aims. The explanatory movement is from the outside in, one could say, rather than from the inside out: instead of starting with what is said, and looking for what the words could mean, the situation where the words fit, is invoked to explain what can be (and is actually being) said (Mey 2001: 542).

I suppose that a great deal can be said about this excerpt. For the time being, I will simply say it reminds me of a film in which the frightened heroine was trying to shout and say something, but was prevented from doing that, her voice never came out. How can we know what she was desperately trying to say, but did not say? The recovery of her intentions, in this case, as in many other similar cases, proceeds from the outside and the movement is from the outside in. We can reconstruct what is going on in her mind only because we can understand what kind of situation she is in.

The considerations above are relevant to my remarks on indirect reports for the following reasons.

The situation of utterance plays a major role in shaping the obligations of the reporter and the degree of accuracy with which the original speech act is reported. Sometimes it determines a more fine-grained report, sometimes a less fine-grained report.

There are contexts in which the hearer is able to separate the reporting speaker's and the reported speaker's voices. While in a default context, I propose a default principle saying that the words used in an indirect report should not lead us too far away from the original statement, I am ready to assert that we also need to investigate those contexts in which the hearer uses some practical means for discerning one voice from another, having at his/her disposal contextual clues and large chunks of world knowledge. For example, if a hearer knows that a certain item of vocabulary belongs to the reporting speaker's linguistic habits, but NOT to the reported speaker's linguistic habits, and s/he knows that the speaker may want to use such a piece of world knowledge to leave implicit part of the message, then s/he will be able to apportion a certain word (the usage of a certain word and the responsibilities involved) to the reporting speaker.

Societal considerations are involved in indirect reports because indirect reports are language games in which in reporting that P, the speaker offers two voices: his voice and that of the original speaker. The reporter does not take responsibility for the embedded voice.

Societal considerations are involved in the fact that certain transformations are precluded in indirect reports, in case such transformations have effects on the speech act reported (threatening to alter it in a drastic way). It is not just words or events that are reported, but speech acts. Since speech acts require appropriate contexts, it is important that the indirect report should present the appropriate context for the speech act narrated, by using words that do not distort that original speech act. As can be easily understood, Mey's idea that speech acts are situated explains the fact that one is not (completely) free to replace one word with another in indirect reports, as, after all, words serve to shape the speech act. The indirect report should provide sufficient contextual clues for the recovery of the original speech act.

That societal information goes into utterance interpretation is now being accepted by leading theorists within pragmatics. There is, in fact, pioneering work by Jaszcolt (2005) saying that the final result of semantic and pragmatic interpretation should take into account socio-cultural defaults. I suppose that the contents of this chapter corroborate Jaszcolt's important, albeit general, considerations.

The problems addressed in this chapter are reminiscent of the issue of mixed quotation, raised by Cappelen & Lepore (2007). In that volume, the authors say that the most typical reasons for preferring mixed to direct or indirect quotation are that the reported utterance is too long to quote, but the reporter wants to ensure accuracy on certain key passages; certain passages were particularly well put; perhaps the words used by the original speaker were potentially offensive to an audience and the speaker wants to distance himself from them; the expression being mixed quoted is ungrammatical or a solecism and the speaker doesn't feel responsible, etc.

The authors opt for a semantic theory of mixed quotation, on the basis of data drawn from the written language. In my review of the book (Capone 2009c), I insisted that by focusing on the written dimension of language, where devices for signaling mixed quotation are quite explicit and of a grammatical nature, Cappelen & Lepore are led astray, as the task of distinguishing quoted from not quoted items is basically a pragmatic one. Fundamentally, in the same way in which mixed quotation poses the pragmatic task of recovering which item belongs to which voice, in indirect reports

more generally there is the problem of telling voices apart. Cappelen & Lepore's problem and my problem intersect, even if, unlike them, I am willing to resort to a theory of pragmemes to show how, in practice, the two voices are to be distinguished.

If we consider indirect reports language games we can see them in a different light from the way philosophers look at them. Indirect reports require a reporter, a piece of language behavior to report and a situation that motivates the reporting and ends up constraining the form of the report. We shall look at indirect reports in the way suggested by Mey, as situated activities, in which the purpose, the participants and the societal rules play major roles in interpretation. As Mey says,

Speech acts, in order to be effective, have to be situated. That is to say, they both rely on, and actively create, the situation in which they are realized. Thus, a situated speech act comes close to what has been called a speech event in ethnographic and anthropological studies (Bauman and Sherzer, 1974): speech as centered on an institutionalized social activity of a certain kind, such as teaching, visiting a doctor's office, participating in a tea-ceremony, and so on. In all such activities, speech is, in a way, prescribed: only certain utterances can be expected and will thus be acceptable; conversely, the participants in the situation, by their acceptance of their own and others' utterances, establish and reaffirm the social situation in which the utterances are uttered and in which they find themselves as utterers (Mey 2001, 219).

From this, we can infer the following:

Indirect speech acts are situated activities. The way the speech act is produced must conform to the rules that pertain to that situation of use. Conforming to the practice of indirect reports implies being competent members of a community of individuals, equipped with communicative competence. This competence has evolved through a mixture of exposure to communicative events and the interplay of cognitive principles.

Indirect reports show a double level of embedding. There is the original act, that was embedded in the situation in which it was uttered; there is the indirect report, which is sensitive to the situation in which it is uttered. Some of the constraints posed by the former context may be overridden by the new constraints of the latter context (for

example an utterance which was inappropriate, misplaced in the former context may be appropriate in the latter context), but in general it can be said that the former context and the interpretations licensed in it will constrain the latter context, UNLESS there are reasons for deviating from such a practice.

Understanding the context in which the original statement reported was made is a step towards a radical theory of pragmemes. In this chapter, I will not discuss in detail a radical treatment of the pragmeme of indirect report, but I must at least mention that the research by Tannen (1989) on talking voices paves the way for an understanding of indirect reports in context. An indirect report is felicitous if it takes into account the original context of utterance, the motives and contingent actions occurring there. Thus it should not be taken for granted that a report that confines itself to repeating certain words is maximally faithful, as it may obscure the contingent function of those words, as well as the speaker's intentions. Radical though that research might be, it does corroborate what I have been saying so far, that is to say that the context of utterance of the original statement constrains the (felicitous making of) the indirect report.

2. A unifying framework: Relevance Theory.

I want to sketch, without going into great detail, the unifying assumptions of this chapter. Many of my considerations given in terms of Principles in this chapter can be connected with the basic assumptions of Relevance Theory by Sperber & Wilson (1986) and Wilson & Sperber (2004). Relevance Theory develops an alternative to the code model of communication, mainly an inferential model in which a communicator provides evidence of her intention to convey a certain meaning, which is inferred by the audience on the basis of the evidence provided. According to the authors, the decoded linguistic meaning is just one of the inputs to a non-demonstrative process which yields an interpretation of the speaker's meaning. The pragmatic process of interpretation is constrained by the following principle:

Cognitive principle of Relevance

Human cognition tends to be geared to the maximization of relevance.

According to RT, an input (whether linguistic or not) is relevant to an individual when it connects with background information to yield conclusions that matter to

him. A positive cognitive effect is a worthwhile difference to the individual's representation of the world.

The authors accept the following:

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.

Pragmatic processing is constrained by the following Presumption:

Presumption of optimal relevance

- (a). The ostensive stimulus is relevant enough to be worth the audience's processing effort;
- (b). It is the most relevant one compatible with the communicator's abilities and preferences.

According to the authors, the hearer/interpreter should follow a path of least effort in computing cognitive effects; test interpretative hypotheses in order of accessibility; and stop when her expectations of relevance are satisfied.

Many authors have written on Relevance Theory, most notably Carston (2002), Bezuidenhout (1997), Rouchota (1992), Powell (2001), and Capone (2008a) (from a more philosophical perspective, which is what I am primarily interested in here) – the main idea that emerges from these publications is that pragmatics provides full propositions out of the fragmentary ones provided by semantics. However, in this

chapter I will mainly address the issue of indirect reports and I will presuppose that semantics and pragmatics work in tandem.

3. Indirect reports as language games.

Wittgenstein writes:

Here the term “language *game*” is meant to bring into prominence the fact that the *speaking* of language is part of an activity, or a form of life (Wittgenstein 1953, 11).

Wittgenstein continues, making a list of language games, and includes among them the game of ‘reporting an event’. It is a natural extension of Wittgenstein’s ideas to say that making an indirect report is a language game – although we must explain why making an indirect report is a more specific language game than making a report.

If we consider indirect reports *language games* we can perhaps put them in a different light from the one which is usually cast on them by scholars in the philosophy of language. Indirect reports require a reporter, a piece of language behaviour to report, a situation that motivates the reporting (one that includes a goal) and that ends up constraining the form of the report (especially of the NPs in the that-clause). The reporter usually does the reporting for the benefit of the hearer (however, it is not difficult to imagine perlocutionary acts such as ‘scaring’ the hearer or putting him off from carrying out an action). This presupposes asymmetrical knowledge between the reporter and the hearer. The speaker would not do the reporting if the hearer were informed of the reported speech (but of course there is the marginal chance that an indirect report has the perlocutionary effect of making the hearer notice that a speaker is a liar – in this case the hearer knows the fact reported, but the indirect report is proffered all the same to focus on the issue of the lie). The reporter qualifies himself and the speaker of the event reported as **samesayers** (to use Davidson’s (1968) words) with respect to the content of the report (the that-clause in “He said that P”). ‘Samesayers’ does not mean that the same words are used in the report and in the speech to be reported: it just means that the report and the speech to be reported have some broad content in common. As Burge says: “To use indirect discourse, one must master the practice of samesaying.

One must be able to use utterances that are relevantly synonymous with the utterances of the original speaker” (Burge 1986, 192).

The practice of indirect reports rests on the following rule:

The reporter X will report what was said by Y (Y usually distinct from X, but sometimes coinciding with X) by using a predicate such as ‘say’ that makes reference to a verbal event of some kind (an utterance), by applying it to Y (the participant whose speech is being reported) and by letting the direct object of the predicate ‘say’ express the **content** of the utterance by Y (at some prior time) by way of paraphrase, that is by letting the that-clause refer to the same situation or event **e** that the utterance **u** reported was actually used (by Y) to refer to, without necessarily using the modes of presentation (of objects and participants) which were actually used by Y in **u**, and in fact allowing context to play a pervasive role in making reference to objects, activities, and participants thereof. In case doubts arise about interpretation of a constituent of what was said or about authorship, quote that part of the utterance, ensuring that the hearer grasps that what follows ‘that’ (quotation marks excised) is the content of what was said, while the quoted part same-tokens an expression actually used.

Such a rule accommodates the problems noted by Davidson concerning reports of utterances in languages other than the reporter’s. The rule above relies on the premise that two expressions are pragmatically equivalent if they express the same content (Jaszczolt 2003) and is reminiscent of Soames’s (1989) position, which has been criticised by Cappelen & Lepore (1997). Of course, this basic rule is incomplete, and I shall attempt to come to better versions of it while I consider crucial examples. A first shot at completing the considerations above is to introduce a rule of use, sensitive to the pragmatic “requirement that the reporter be maximally faithful to the words of the agent unless there is a reason to deviate” (Soames 1988, 123) (see also analogous considerations for belief reports in Salmon’s work (Salmon 1986), and Saul 1998 for a critique).

A second shot is that contextual considerations (e.g. the formality/informality of the situation) can increase the need for a more fine-grained report, or coarse-grained

report. In other words, a speaker has to fine-grain an indirect report in order to adjust to the situation of use.

Making an indirect report is a language game that is more specific than making an assertion: in asserting P a speaker merely offers her own voice (unless ironies are involved; see Giora (2003) for an original and important view of ironies); in reporting that P, the speaker offers **two voices**: her voice and that of the speaker in the original speech event (see also Vološinov 1973) . The reporter does not take responsibility for the embedded voice (except in so far as he attributes it to a speaker or another). Reported speech is usually elicited (A: What did John say?) or prompted by the desire to offer H (the hearer) an essential clue to a solution of a problem viz. some piece of information that is contained in (or is a consequence of) the speech reported (I admit I was influenced by Devitt's 1996 pragmatic approach to belief reports, which I freely extend to indirect reports). Reporting speech is a language game because it involves some principles. It is a principle-based activity. Furthermore, the language game of reporting (someone's) speech is sensitive to contextual factors and the context of speech determines whether a report should be more (or less) fine-grained.

The problem as I see it is that reporting speech (indirect reports) is a language-game of its own that is severely constrained by the fact that it displays **two voices**, the reporter's and the original speaker's, and there is a **tension** between them, such that the interests and the point of view of the reporter cannot prevail over those of the original speaker. Furthermore, the specificity of the language game 'reporting (one's) speech' consists in the tension between the two voices and in the fact that none of them prevails over the other. The specificity of the language game 'reporting speech' also consists in the fact that the report partially answers the question "would I accept the report (the way in which the report has been put), were I the original speaker whose speech is being reported?". This (salutary) question serves to eliminate possible distortions of what was said. Indirect speech reports represent what Bakhtin calls 'discourse of the third type', that is discourse which does not just express the speaker's voice or that of a third person, but multiple voices (Robinson 2003, 107; Bakhtin 1984, 187).

The situation of utterance plays a major role in shaping the obligations of the reporter and the degree of accuracy with which the original speech is reported. Sometimes it determines a more fine-grained report, sometimes a less fine-

grained report. We can extend somewhat the principle we came up with by incorporating the notion of the specificity of the language game we are dealing with:

Paraphrasis Principle¹

The that-clause embedded in the verb ‘say’ is a paraphrasis of what Y said that meets the following constraint: should Y hear what X said he (Y) had said, he would not take issue with it, but would approve of it as a **fair paraphrasis** of her original utterance.

This is somewhat reminiscent of Burge’s position (Burge 1986) that:

The point of indirect discourse might be fairly taken to be to introduce and produce a given utterance that gives the content of the original speaker’s utterance (Burge 1986, 196).

Now it should be clear why the language game ‘indirect report’ is more specific than that of ‘assertion’. In factual assertions, only the speaker and the hearer are involved; the speaker takes responsibility for what he takes to be the case in the actual world. In the case of indirect reports, we have got a speaker’, a hearer and a speaker’’ (the original speaker), and both the speaker’ and the speaker’’ take responsibility for the content of the that-clause in case all goes well, that is to say, in case the indirect report is felicitous.

There is another aspect to take into account in our argument that indirect reports are specific language games. Matters of form also go into an evaluation of whether the reporter made a correct report. In fact, as Dummett (2003, 110) says, one may assent to a statement without being prepared to make that statement (since it may be objectionable in other ways, e.g. by being insulting).

Thus we also need the Paraphrasis/Form principle for indirect reports:

Paraphrasis/Form Principle

¹ This position is somewhat reminiscent of Seymour’s (1994) treatment of indirect reports, in which reference to a translation of the reported sentence is explicitly incorporated in the semantics of indirect reports.

The that-clause embedded in the verb ‘say’ is a paraphrase of what Y said, and meets the following constraints: should Y hear what X said he (Y) had said, he would not take issue with it, as to content, but would approve of it as a fair paraphrase of his original utterance. Furthermore, he would not object to vocalizing the assertion made out of the words following the complementizer ‘that’ on account of its form/style.

Now, it should be clear that the way I have formulated this principle is that of a generalization or a rule. It would be good, however, if this were not an isolated rule, but if it could be reducible to a more general cognitive principle (or if it could be connected with such a principle). The Form/Paraphrase Principle could be explained by resorting to Sperber & Wilson’s (1986) theory of Relevance. According to Relevance Theory, an assumption is more relevant if it has greater cognitive effects or if it has smaller processing efforts. An indirect report in which it is the reported speaker’s voice and not the reporting speaker’s voice that prevails is a case in which the hearer of the report is put to smaller processing costs. Thus it follows that the indirect report will merely serve to give voice to the reported speaker (unless otherwise indicated, by contextual clues).

Of course, if the discourse is rich enough and offers contextual assumptions that help separate the reported speaker’s voice from the reporting speaker’s voice, the greater processing costs involved in the report will be offset by greater contextual effect.

A very interesting objection could be levelled to the Paraphrase/Form Principle:

Suppose that I mistakenly think that Sam is a woman. I utter ‘Sam is a philosopher’. You report my utterance by uttering ‘X said that he is a philosopher’ (my own addition: ‘he’ being used in combination with a demonstrative gesture). I would take issue with ‘he is a philosopher’ (or otherwise object to its form), because I think that ‘he’ is not the right grammatical gender for Sam. So, either you cannot replace ‘Sam’ with ‘he’ or the author’s Paraphrase/Form Principle – according to which the ‘that’-clause in your report of my utterance must be such that, “should [I] hear what [you] said [I] had said, [I] would not take issue with it and would not “object” to it on account of its form/style” – is false.

The example is rather tricky. However one can hope to deal with it as follows. Surely, if there was no way to settle the issue of Sam's sex, the Paraphrasis/Form Principle would be a serious problem. However, there are ways to surmount the problem provided that there are publically available and objective criteria for establishing sex (usually by exterior physical appearance). Suppose X utters 'Sam is a philosopher' and I report 'X said that he is a philosopher'; surely X can object to my report because he is falsely persuaded that Sam is a woman and until he is so persuaded he will object to my report. This is not to say that he is entitled to object to my report. All we need is a notion of being entitled to object to the report. All we need is a principle of use to the effect that the report is felicitous provided that the reported speaker is not entitled to object to it. The example is also tricky because it makes use of the English language, where names can be ambiguous as to gender. But in many languages of the world there are no such ambiguities, and, thus, the reviewer's story would not work there.

Another line of objection could be the following:

Depending on the context, I needn't be beholden to the original speaker's 'approval' of my paraphrasis as fair, nor need I avoid manners of speech which the original speaker would shy away from. In such contexts, if John said of a person x that he will be coming to the party, my report to that effect is true whether I refer to person x politely, as John would approve of, or impolitely, as (let us imagine) my hearer would approve of. John may, upon hearing my report, demur: "Well, I don't know why you'd call x a jerk but, yes, I did say he was coming to the party". The Paraphrasis Principle and the author's other remarks are intended to rule out contexts of indirect reporting that seem to allow this type of license with the original speaker's words.

There are two points to be noted, in this objection. The hypothetical reply above by someone whose speech was reported by a mixture of the reported speaker's voice and of the reporting speaker's voice at least signals a complaint, and a complaint is something which one voices when a trouble has been noted. Furthermore, one has to ask whether the reported speaker who said "Well, I don't know why you'd call x a jerk but, yes, I did say he was coming to the party" would also be inclined to say

“You reported what I say all right” – it is not clear that s/he would also exhibit such an inclination. The second point to be noted is that the objection starts with ‘Depending on the context’ – and thus I am inclined to grant that there are contexts in which the hearer is able to separate the reporting speaker’s and the reported speaker’s voice and in such contexts the Paraphrasis/Form Principle will be blocked. The versatility of such principles seems to attest to the fact that we are dealing with procedures determined by language use, which are therefore sensitive to language use. We are dealing with the social practice of indirect reports, which rests on usage and on regularities of usage. If a context is such that a hearer can separate the reported speaker’s from the reporting speaker’s voice, the Paraphrasis/Form Principle is no longer operative. It is operative only in those contexts in which it is not easy to keep the voices apart. I am conscious that the discussion at this point become very interesting and stimulating, but for the sake of time I have to stop here.

4. The logic and structure of indirect reports

I want to outline briefly the logic of indirect reports. Indirect speech reports are micro-narrations: in particular, they narrate certain events (deemed to be of interest to the hearer) that occurred at some time prior to the speech report and that amounted to saying something with appropriate intentions (here intended *à la* Habermas 2001, as teleological entities; see also Peirce 1958, 414)). Assertive speech acts are events reportable in indirect speech reports. But now, in the case of indirect speech reports, there are two speech acts: the original assertion (subsequently reported) and the micro-narration of the original assertion (another assertion). It is reasonable to assume that the original assertion was connected/related to a certain situation (a complex of cognitive states, goals, desires) and had a bearing on the formation of certain decisions on the part of some of the participants to the speech event (in this way, cognitive effects, to use terminology by Sperber & Wilson 1995, are maximal; see also Wilson 2000; Capone 2001). The indirect report, analogously, relates to some (new) situation (a complex of cognitive states, goals, desires, etc.) and since, following Wilson 2000 and Kamp (1990, 70), it leads to an inference concerning the beliefs of the original speaker, it is tantamount to a belief attribution; thus, following Kamp (1990, 30), it ends up having a bearing on the formation of decisions on the part of some of the participants (to the speech event), in the sense that if the hearer is

informed of a certain situation that has a bearing on taking a decision, then he will be willing to take it given this piece of information. The way the indirect speech report can bear on a certain decision is to propose what another person said (asserted) as a **source of knowledge**. If the original speaker qualified as a reliable informer, then what he said can be counted on for the formation of appropriate beliefs that have a bearing on the formation of current decisions relating to the current complex of cognitive states, goals, desires. Alternatively, what the original speaker said can be subjected to further inquiry and can be contrasted with conflicting pieces of knowledge. So far, I have presupposed (rather generously) that there is (normally, albeit not always) an inferential transition from what a speaker A says to what she believes. The way I see the difference between a factual assertion and an indirect report is that the factual assertion is merely a response to the interests of the hearer at *t* (the time of utterance); instead, an indirect report embodies a sensitivity to two situations, *s* and *s'*. *s* is the situation in which a factual assertion (in any case the original speech act reported) was uttered in response to the interests of recipients *R*. *s'* exhibits a sensitivity to both recipients *R* and *R'*. It goes without saying that if modes of presentation reflect the interests of the hearers (as well as of the speakers), then an intersection of the interests of *R* and *R'* should be taken into account in the choice of mode of presentation (as Sperber 1996 says, communication normally slightly transforms the message). As should now be obvious, the methodological approach to be used in this chapter is to accept that since users and their language are at the core of all things pragmatic, the world of users is the very condition for doing pragmatics (Mey 2001, 29; Haberland and Mey 1977, 1-16).

5. Restrictions on transformations.

Suppose that a certain referent, say *A*, features in both the original speech act and in the reported (indirect) speech act. Suppose that in the original speech act, *A* is presented through a mode of presentation *M*, whereas in the indirect reported speech act, it is presented through a mode of presentation *M'* and also suppose that *M* and *M'* happen to be distinct. Surely, the context has to determine the level of detail associated with the mode of presentation and whether a more **fine-grained** linguistic expression is to be preferred to a more **coarse-grained** expression. Yet, the latitude of the choice can be restricted a priori. Consider the following example from Higginbotham's lecture notes (2004/2005):

(1) Galileo said that the earth moves.

(2) Galileo said that the planet in which Arnold Schwarzenegger is a governor moves.

Higginbotham says that (2) is “ridiculous” (even if it is obtained by substituting a coreferential expression for ‘the earth’). He does not explain, though, why he considers examples such as (2) (obtained through substitution of identicals from (1)) ridiculous – presumably his explanation is that it was unlikely for Galileo to have any thoughts about Schwarzenegger and, thus, to use a mode of presentation making reference to such thoughts. However, I believe the explanation cannot be that verbs such as ‘said’ block Leibniz’s Law, because indirect speech reports that are not pedantic do involve some latitude in the choice of mode of presentation for a referent presented via a different mode of presentation in the original speech act.

There is a reason why we are barred from shifting freely from (1) to (2), a reason that does not consist in the blocking of Leibniz’s law (a statement will result in a coextensive statement if an NP is replaced with a coextensive NP; see also Jaszczolt 1999) in connection with verbs of propositional attitudes. Suppose that we have some latitude in shifting freely from one mode of presentation to another in paraphrasing a sentence (subject to the constraints formalized above), thus making the application of Leibniz’s law tolerable. Although we are allowed freedom in choosing modes of presentation of the same referent, we are barred from choosing a mode of presentation which, once inserted in the proposition said, would amount to our accepting that the original speaker said something implausible. Surely Galileo could not have said that the planet in which Arnold Schwarzenegger is a governor moves, because Arnold Schwarzenegger was not living at Galileo’s time.

A semantic story could, however, compete with a pragmatic story, in this particular example. It could be said that the contrast between (1) and (2) could be handled in terms of scope phenomena. Presumably, such a story assumes that ‘say’ is, in respect to quantificational structure, similar to ‘believe’ and allows the existential quantifier to take wide scope. Semantic theory allows two logical interpretations of (2): one true and one false. Yet, as I maintain, we almost always interpret (2) in the second way. In the case of (2), the two logical accounts are:

(2a) There was an x, x is the planet where Arnold Schwarzenegger is a governor: Galileo said that it moves;

(2b) Galileo said that: There was an x, x is the planet where Arnold Schwarzenegger is governor and x moves.

The pragmatic point is that in normal circumstances an utterance of (2) is to be interpreted as (2b), rather than (2a). (2a), though not doubt rare, may not be impossible, for instance if it is a conversation in which the participants know that (1) is true, but don't know what Schwarzenegger is governor. Even then (2) would probably be odd, and it is pragmatics that establishes this. Of course the problem, for pragmatics, is to specify why logical form (2b) gets chosen. Suppose we stick to the notion I have so far expressed that an indirect report can express one or more voices. (2a) is clearly the case where there are two voices to be taken into consideration: the reporting speaker and the reported speaker. (2b) is clearly the case where there is only one voice to be taken into consideration: the reported speaker's. If we adopt Sperber & Wilson's (1986) idea that an assumption (an input) is more relevant if either (a) it has greater positive cognitive effects or (b) it has smaller processing efforts, we can easily understand why (2b) gets standardly chosen. In fact, (2b) involves smaller processing effects, as the hearer must simply allocate the reported speaker's voice, while it does not need to separate it from the reporting speaker's voice.

There are many examples, similar to Higginbotham's case, that need to be accounted for. In fact, while in a few cases replacement of an NP with another co-referring one is barred, in general, unless problems arise, such substitutions are possible in indirect reports, especially when they aid **understanding on the part of the hearer**. So the real problem, for the semanticist, is not to say that substitutions are not possible in indirect reports, but to specify those cases in which substitutions are blocked, explaining why they are.

One such case is the following. Suppose John utters: "A fortnight is a two week period". Assuming that in transformations from direct to indirect speech (NP) substitutions are licit, one could then report the following: "John said that a fortnight is a fortnight". This is a tautology which did not appear in the original speech act. An

easy way out of this problem is to say that if no **tautology** appeared in the original speech act, it must not appear in the reported speech act, because it will give the hearer the impression that it was also present in the original speech act (which is not the case).

But this is only an easy way out. The real problem is whether “John said that a fortnight is a fortnight” does capture the original speech act. There is no reason to believe that it does. The original speech act, in fact, is motivated by the goal of **explaining** one word by making recourse to an easier word. The reported speech act, instead, does not explain anything, because it just establishes an identity between a word and itself. Since the illocutionary force of the original speech act (**explanation**) is lost due to a transformation that amounted to replacing a term with a co-referential one, it is fair to say that such substitutions are not legitimate in the transition from direct speech to indirect speech reports. This presumably follows from a general semantic principle, worked out by Alston:

A meaning of a sentence fits it to play a distinctive role in communication just because that meaning consists in the sentence’s being usable to perform illocutionary acts of a certain type (Alston 2000, 282).

We are now in a position to explain what is odd with substitutions of terms having the same sense within indirect reports. Platts (1997) is puzzled by the fact that, contrary to what might be expected given the acceptance of Frege’s explanation of failure of substitution in intensional contexts (the view that in intensional contexts NPs refer to their senses), in indirect reports one obtains strange results by substituting an NP with a distinct one which has the same sense (so, not only coreference, but sense identity are at stake here). For example, see what happens when we replace ‘oculist’ with ‘eye-doctor’ in the sentence: “John said that an oculist is an eye-doctor”. The result is the sentence “John said that an eye-doctor is an eye-doctor”, which may be rejected as false (or pretty inaccurate) by John, since his original speech act did not contain a **repetition** unlike the sentence which is the result of the substitution. A way out of the puzzle for Platts is to deny that ‘oculist’ and ‘eye-doctor’ have the same sense (but this would amount to denying that the words are synonymous). Another way out of the puzzle is to say, as I have done for the previous example, that substitutions of co-referring or synonymous expressions

are prohibited when they severely distort the nature of a speech act originally proffered by the speaker. In the case we are discussing, the original speech act is an **explanation**, in which a more difficult word is explained via a less difficult one. If we replace the more difficult word with a synonymous less difficult one, we lose the pragmatic force of the utterance, which was meant to be an explanation.

We can also frame this explanation in the framework of Relevance Theory. If a transformation results in a different speech act being understood, then it is clear that what bans the transformation is the Relevance theory assumption that greater processing efforts (as the ones involved here) will result in a less relevant proposition. A transformation that obscures the original speech act will result in a blend of two voices, whereas an interpretation that merely promotes one voice keeps processing efforts lower.

6. Indirect reports, indexicals, and speech acts.

It is good at this point to dwell on the interaction between indirect speech reports and indexicals. Among modes of presentation, the indexical has a special status. When a person says “I”, she is presenting herself not only as the person who is speaking, but also as the experiencer of a special state of perception: “I” is a mode of presentation that presents the subject of a special type of experience: a person who has privileged access (direct access) to one’s thoughts and being (see also Zeevat 1997, 163). There is a substantial difference between saying:

(3) I am happy;

(4) Alessandro is happy.

(3), through the mode of presentation “I”, makes it clear that the state of perception “happy” (attributed to the subject) is experienced directly by the person who speaks or thinks; (4), instead, attributes a state to a subject by implying that such an attribution is based on an external mode of inference:

e.g.

Alessandro is dancing at a party;

Alessandro dances only when he is happy

∴

Alessandro is happy.

It would be misleading to utter the third person statement (4) in case the person who speaks (or thinks) is speaking of himself, as (4) implies that the attribution is based on an external mode of inference, whereas an internal mode of inference (direct perception) is implied in first person attributions.

Levinson (2004) notices that it will not do, in many cases, to replace the indexical pronominal “I” with a third person paraphrasis (e.g. Stephen Levinson). The contexts he has in mind are constituted mainly by identity statements such as:

(5) I am Stephen Levinson.

It is obvious that, by replacing the indexical pronominal “I” with a third person paraphrasis, one obtains a (relatively) uninformative sentence:

(6) Stephen Levinson is Stephen Levinson.

Now, while Levinson’s point is to show that the indexical mode of presentation has a special status, this example throws light on the mechanism of indirect speech reports. In fact, it goes without saying that it will not do to report the utterance (5) by saying:

(7) He said that Stephen Levinson is Stephen Levinson.

Of course, it should be explained that ‘He said that I am Stephen Levinson’ and ‘He said that Stephen Levinson is Stephen Levinson’ do not have the same meaning (semantic import), as in the former sentence we find a pronominal while in the same syntactic position we find a full name in the second sentence². Even if the two sentences in context may end up having the same truth-conditions, intuitively we know that they count as different assertions. So it appears that a semantic story plus saturation of indexical elements does not suffice to distinguish between the two resulting statements. We need to know the purpose with which the speech act connects.

² I thank a referee for this not uninteresting observation.

In fact, the indirect speech report should also inform the hearer of the type of speech act the original speaker made, in saying what he said. (7) is not a good indirect report because, by replacing the indexical mode of presentation with a third person paraphrase, the speaker has managed to obscure the original relationship between the indexical mode of presentation and the speech act actually made: the speaker was **introducing himself** by an identity statement. In introductions, the indexical mode of presentation is crucial, because it renders the identity statement informative. By replacing the indexical mode of presentation with a third person paraphrase, the identity statement is made *ipso facto* uninformative as in (8):

(8) Stephen Levinson is Stephen Levinson.

While this is still an identity statement, it is not an introductory speech act. The introductory speech act needs to establish an identity between the indexical mode of presentation and another mode of presentation. This is the change it has to bring about in the existing state of affairs, to put it with Mey's (1993, 111) words. The consideration that pragmatic equivalence (or the lack of it) explains why substitution of an NP is sometimes not licit certainly follows from Jaszczolt's (2003) remark that equivalence of meaning is not just of the semantic type, but also of the pragmatic type (sometimes, according to Jaszczolt two completely different sentences are pragmatically equivalent; her examples are based on contrasts between languages).

So a plausible generalization for indirect speech reports is that:

a mode of presentation cannot be replaced with a third person paraphrase if such a change will inevitably obscure the kind of speech act which was made in uttering the original sentence.

Again, I would like to connect my considerations with Sperber & Wilson's (1986) Principle of Relevance. If an indirect report effects a transformation that has a bearing on the resulting speech act, it will put the hearer to greater processing costs. Suppose we can freely transform from 'He said he is Stephen Levinson' to 'He said Stephen Levinson is Stephen Levinson'; we move from an indirect report from which we can evince the speech act performed by the original speaker to an indirect

report where the hearer must use greater inferential powers to infer the original speech act. The hearer has clearly to embark on greater processing costs and the input is thus not ideally relevant.

7. Choice of mode of presentation and the hearer.

It may be thought that the choice of a certain lexeme (instead of a competitor) in an indirect speech report is not a matter of having respect for what the original speaker actually said, but a matter of ensuring that the hearer be able to identify the referent. If the purpose of an indirect speech report is to inform a hearer H of what a speaker S said in a prior speech act, so that H can utilize the knowledge imparted by what S said (assuming that what he said transmitted S's knowledge), then it goes without saying that it would be better to choose an NP (or mode of presentation) recognizable by the hearer (otherwise the purpose of the indirect speech act will not be fulfilled). Habermas is quite right in saying that

By describing behaviour as an intentional action, we take the perspective of the actor himself; but this agent's point of view signifies a two-tiered intentional relation to something in the world, namely the relation to the cognitive representation of reality that is valid for the agent and to the subjective attitude that the agent takes towards this representation of reality (Habermas 2001, 113).

The fact that indirect reports can use different modes of presentation from those used in the original speech acts, shows that Habermas' two levels of intentionality are at work in indirect reports³: the Russellian proposition of the that-clause (the cognitive representation of reality Habermas alludes to) and the intention to use a mode of presentation that is more familiar to the hearer so that the hearer can utilize the knowledge imparted with the that-clause to carry out whatever action he deems necessary. The subjective attitude Habermas speaks of is the intention to act in a certain way (propositional attitudes such as desiring to act in a certain way).

Indirect reports, from what we have seen, are often associated with perlocutionary purposes; thus, it is correct to say that "minds are partial causes of events in the

³ It is clear from the discussion that Habermas aims to reconcile Brentano's notion of intentionality (thoughts are intentional in that they are directed towards objects and contents) with a teleological notion of intentionality.

world and in other minds” (Zeevat 1997, 156). It is uncontroversial that when the indirect report has got the purpose of letting a hearer utilize what an original speaker said so that the hearer of the speech report can form some knowledge of a situation, then an NP that can guide the hearer towards identifying the referent must be chosen. However, there are constraints on such a choice. The NP chosen **must not** distort what the original speaker said in such a way that, faced with the indirect speech report, the original speaker is likely to say that he does not recognize his intention in that report.

8. Indirect reports and expressives

As an authoritative linguist has said, we can sometimes be faced with a text or an utterance “whose ambiguous linguistic form makes it ‘double-voiced’...”with an ambivalence of voice” (Fairclough 1992, 108). The practice of indirect reports involves being able to separate what is attributable to the original sayer and what is attributable to the current speaker, even if both things appear in the that-clause. So a good Principle is the following:

Do not take everything that appears in the that-clause of the indirect report as belonging to the voice of the original speaker whose speech act is being reported.

A complementary Principle is the following:

Separate the elements of the that-clause that contribute to the voice of the original speaker from those that contribute to the voice of the reporter by exploiting contextual clues – assuming that these are sufficient for the purpose of separating the two voices.

An important example that illustrates the role contextual clues play in separating the two voices is the following drawn from Potts’s (2005) discussion of conventional implicatures and, in particular, of expressive acts (the author takes expressive acts to be a special class of adjectives that can never contribute to the at-issue content):

9)

Edna is at her friend Chuck's house. Chuck tells her that he thinks all his red vases are ugly. He approves only of his blue ones. He tells Edna that she can take one of his red vases. Edna thinks the red vases are lovely, selects one and returns home to tell her housemate:

'Chuck said I could have one of his lovely vases!' (Potts 2005, 18).

Pott says:

"We easily recognize that Edna is contributing the adjective (lovely); the utterance expresses two propositions: (i) that Chuck said Edna could have one of his vases; (ii) Edna thinks Chuck's vases are lovely (Potts 2005, 18).

It is not clear that Potts is completely right about what he says. Surely Edna's housemate need not understand that 'lovely' refers to Edna's and not to Chuck's voice in the absence of appropriate information. In fact, he has got no contextual clues to separate Edna's from Chuck's voice in the that-clause attributed to Chuck. Perhaps he can do so only after Edna narrates the whole story, as we have it above. Instead, Edna and the readers, who have access to the whole story, are able to separate Edna's from Chuck's voice. One obvious way to do so is to avoid attributing to Chuck elements of the that-clause in case it is clear from the context that they do not reflect (they contradict) Chuck's opinion. In this case, context is used to filter out from the that-clause elements (of meaning) that do not reflect the opinion of the original utterer. Without full contextual knowledge, it is impossible to separate the original speaker's from the reporter's voice (some underdetermination of meaning will result).

As usual we may want to frame this discussion within the context of Relevance Theory. If one accepts that an assumption is less relevant in case it puts the hearer to undue processing efforts, then one at the same time explains why, in the case of lack of contextual information about Chuck's beliefs, Edna's utterance 'Chuck said I could have one of his lovely vases' will be typically interpreted as expressing Chuck's and not Edna's voice. Having two voices in the same utterance would make processing costs greater, and there would be no contextual assumptions to yield a plausible interpretation on how to distinguish the two voices. On the contrary, when

there are contextual assumptions that allow us to distinguish the two voices, the extra processing costs will be offset by rich positive contextual effects.

It could be said that it is the semantical assumptions that makes it the case that there is a way of interpreting (9) which makes what Edna says true, in case we know that Chuck thinks the vases are ugly but Edna does not. Presumably it could be thought that we might analyse (9) semantically as characterised by a scope ambiguity and as being associated with two logical forms:

(9a) There are two lovely vases such that Chuck said I could have one of them.

(9b) Chuck said that there are two lovely vases and I could have one of them.

As in the previous case of ambiguous logical forms, pragmatics will intervene in promoting one, rather than the other logical form. As I said previously, the principle of Relevance will give expectations that are precise and predictable enough to guide the hearer towards the speaker's meaning. Since (9a) involves greater processing efforts, (9b) will be selected instead, unless there is a rich context in which through contextual assumptions one will be able to separate the two voices (the reported speaker's and the reporting speaker's) and thus (9a) could be promoted as a reasonable interpretation.

9. Final remarks

This chapter endeavoured to contribute to the theory of indirect reports, by the notion of linguistic practices (certainly a societal notion) and by the formulation of a number of principles that are central to these practices. Yet, it would be good if we could find out some general rationale for the fact that sometimes co-referential expressions are inter-substitutable in that-clauses (of indirect reports), while sometimes they are not. I believe that a good filter is to be found in speech act theory. If we agree that saying is not just a locutionary, but also an illocutionary act, then transformations of what a speaker originally said can be tolerated **provided that** the illocutionary act originally made is preserved in the final indirect speech report. This is the way the filter works. Now we can understand why an NP cannot be replaced with an NP that contains an epithet (e.g. "That bastard"), since, by so doing, the assertion will *ipso facto* be transformed into an **insult**, which is a different speech

act from an assertion. This is not to say that one cannot disagree on this. I expect the following objection:

But if you utter 'I voted for Obama' and I am a McCain supporter, when I am talking to my fellow McCain supporters I can report your utterance by uttering 'X said that X voted for that bastard'. This use is called 'pseudo de re'.

However, there is in this alleged counterexample a deictic element combined with the epithet (which correctly deserves the name of 'pseudo de re') that gives the hearer enough clues to tell apart what is said by the reported speaker and what is said by the reporting speaker. Thus, it is not a good counterexample. A good counterexample would be something like the following: 'Obama said that he voted for McCain the bastard'. Here there are no sufficient contextual clues to tell apart what Obama said from what the reporter said, and thus the indirect report is not felicitous.

However, suppose one takes the objection above seriously. Suppose that the grounds for not accepting my story that some transformations strongly imply that the original speaker is responsible for the epithet are that there is no reason to think that contexts in which the hearer of the indirect report knows that the reporter has a negative attitude toward the individual and knows that the reported speaker does not are rare. In such contexts there would be no tendency of the hearer to impute the epithet to the reported speaker.

I agree that the counterexample discussed before (the Obama / McCain example) is suitable to such contexts and I agree that in such contexts the hearer of the report will use whatever contextual clues are available to distinguish the reporter's from the reported speaker's voice. The problem is what happens in a context where there are no such strong contextual clues. What is the default procedure in such contexts? I suppose that the purpose of an indirect report is to give voice to the reported speaker's voice and it is this that prevails. The alternative I envisage to my own solution is that one tries to establish from case to case, without any default procedure, whose voice is in question. However, a default procedure is probably what captures

best the fact that indirect reports evolved with a precise function in mind: reporting at some level of precision the voice of the reported speaker.

Conclusion

The theory of indirect speech reports will be advanced considerably if we place this issue, which surely pertains to the philosophy of language, in a societal perspective. It is very instructive to investigate the societal principles having a bearing on the transformations involved in the transition from direct speech to indirect speech reports, in particular, those pertaining to modes of presentation. The use of some, rather than some other, mode of presentation is a matter of being situated in a certain context. It follows that contextual information severely limits transformations pertaining to modes of presentation. Such a story become even more plausible if it is framed in the context of Relevance Theory and if one can show that the constraints on the interpretation of indirect reports follow from the principle of Relevance. This is what I have done.

I end this section by voicing a thorny objection to myself. Do I really believe that my societal considerations can be reconciled with Sperber & Wilson's theory of Relevance? Surely there may be more than one answer to this question. My modest answer is that, even if Relevance is a cognitive principle, it could be used to justify social practices. There is no reason why social practices should not conform to a cognitive principle that has explained much of language behaviour so far. Language behaviour is, *ipso facto*, a social practice. On the one hand, social practices can be sensitive to cognitive principles and be influenced by them; on the other hand, since social life has a life of its own, habits and practices which become consolidated are propagated and transmitted to further generations as unmotivated practices. Sperber (1996) shows that the Principle of Relevance plays a major role in creating cultural practices. Cultural practices which conform to the Principle of Relevance have greater chances of survival and will be more likely to be transmitted from one generation to the next. In the case of linguistic practices, I propose that on the one hand, the principle of relevance will play a role in promoting one, rather than another, cultural practice; on the other hand, it will play a role in the preservation and transmission of the practice, as other competing practices will be ruled out due to

cognitive considerations. This is what makes certain practices better candidates for propagation than others.

There is one further sense in which the practices we have considered in this chapter, as determined by the principle of Relevance, are social. The fact that the Principle of Relevance creates expectations that are precise and predictable enough to guide the hearer toward the expected meaning guarantees that what is cognitively convenient also becomes socially convenient, as one will be certain that an interpretative pattern or a social practice that is more economical will be more likely to be adopted by society. There is every reason to believe that society will orient towards those practices that are more in line with the Principle of Relevance.

Chapter 8

Indirect reports as language games.

Abstract

In this chapter I deal with indirect reports in terms of language games. I try to make connections between the theory of language games and the theory of indirect reports, in the light of the issue of clues and cues.

En outre, comme tout autre terme predicative, 'jeu de langage' n'acquiert un sens qu'à travers ses applications, c'est-à-dire, dans le jeu qu'on joue avec lui. (Dascal et al. 1996, 1371).

Introduction

In this chapter I shall investigate the issue of indirect reports within the framework of societal pragmatics (Mey 2001) or socio-pragmatics (Dascal 2003). Indirect reports are language games based on norms or principles which are rigidly enforced; like language games they are practices which need to be learned in environments where actors engage in the practice and where there are people ready to enforce norms.

Indirect reports are based on an interplay of voices. The voice of the reporter must allow hearers to 'reconstruct' the voice of the reported speaker. Ideally, it must be possible to separate the reporter's voice from that of the reported speaker. When we

analyze the language game of indirect reporting, we ideally want to establish which parts belong to the primary voice (the reported speaker's voice) and which parts belong to the reporter's voice. If we have ways to recognize separate styles, separate grammars, etc. then the process of separating voices will be easier. We should however expect that the reporter will do what he can to make it easy for us to distinguish the two (or more) voices. In other words, UNLESS there are clues which can lead the hearer to recognizing separate voices, the reporter should do his best to represent (without interpolations) the reported speaker's voice (This polyphonic approach is indebted to ideas by Bakhtin 1984, 1986).

The literature on indirect reports lives on the heritage of Davidson (1968) and his critics. In this chapter, I will not directly address the issue of whether the semantics proposed by Davidson is correct or not. Despite my neutrality, I claim that Davidson's treatment gives us numerous clues on how to deal with the practice of indirect reporting. Consider the famous report:

1) Galileo said that the earth moves

Davidson's proposal is to treat (1) as if it meant:

Galileo uttered a sentence that meant in his mouth what 'The earth moves' means now in mine.

I know well that this is only part of Davidson's picture; however, this is the part which can be exploited in the discussion of the phenomenon of 'voicing' which this chapter is about. Davidson's treatment makes it clear that there are two voices (the reference to the mouth is clearly a reference to the mouth proffering the speech or the utterance). Davidson, of course, is interested in the equivalence between the speech proffered by the reported speaker and the speech proffered by the reporter, an equivalence being based on intended meaning, not only on sentential meaning. We

are not dealing with sentential meanings because Davidson makes an explicit reference to a proffered sentence, that is to say an utterance.

So we have a solid platform from which to start working. In the following sections I shall discuss the phenomenon of voicing and the problem of how to separate voices.

1. The transformation problem

It would be naïve to believe that an indirect reporter has a duty to report verbatim what was said. If she had such a duty, then we would not be able to make any difference between indirect and direct reports. Instead, we all have a firm grasp of the difference between:

- (2) She said: I am happy;
- (3) She said she was happy.

The difference is not only one based on transformations of deictic elements (e.g. ‘I’ → she; present tense → past tense; first person morphology → third person morphology), but one based on a greater number of transformations. For example it is normally licit to make an indirect report by summing up what a speaker said, rather than reporting all the elements he uttered. (See also Wieland Forthcoming). It is possible to omit adverbs, adjectives, modifiers, in a sentence, without distorting it, if the purpose which the indirect report serves is fulfilled by the omission of certain words. Of course, if the omission of a certain word results in an utterance which somehow distorts the reported speaker’s message, then the omission should not be tolerated. There are small transformations of a message which serve to modify the message considerably. Consider what was done by the blacksmith who created the inscription ‘Arbeit macht frei’ over the Auschwitz main gate. By creating a letter which was upside down, he managed to express his own negative attitude to the message which he was reporting by his manual work.¹ The change of a letter in the

¹ At Auschwitz, the sign was made in 1940 by Polish political prisoners headed by Jan Liwacz (camp number 1010). The upper bowl in the "B" in "ARBEIT" is wider than the lower bowl, appearing to some as upside-down. Allegedly it was made on purpose by political prisoners to make a signal about what was actually happening behind the gates of Auschwitz.

message seems to stress the awareness that he was used to report a message and his moral reluctance to be so used. In other words, he was distinguishing his voice from those of the Nazi – in a way he was a polyphonic ‘animator’ in the sense of Goffman.

Transformations of the original message of indirect reporting include expansion of the original message, like the following case. Prisoners in Auschwitz ridiculed the German cynicism of the slogan by saying *Arbeit Macht Frei durch den Schornstein*. They were clearly reporting a message by expanding it so as to recuperate its (real) intended message. The general point of transformations concerning indirect reports is to make explicit otherwise implicit components of the message. It is licit in general to render explicit what relevance theorists call the ‘explicatures’ of an utterance. In neutral cases, reports that expand the message so as to include explicatures need not express a critical attitude to the message, as was the case in the ‘Arbeit macht frei’ report.

It is less clear that one can expand the message by including elements obtained through logical inference:

The Nazi used the slogan ‘Arbeit macht frei’ in order to indulge in a morbid exaltation of their crimes.

This may not belong to ordinary practice of indirect reporting, but it is clear that this is a practice in which all historians indulge.

A transformation which, instead, belongs to daily practice is the following. One can report an utterance proffered in a different language by someone who spoke (only) that language, through the reporter’s language. So, suppose I report that Cesar said that he came, he saw and he won; it will be implicit that the words used by Caesar were words of Latin, and not of English. My report, therefore, only relates the content of what he said, and does not report it verbatim. The report involves a translation. But then, one may say, translation allows all sorts of transformations, as one clearly sees if one for example reads the many translations of Thomas Grey’s ‘Elegy written in a country churchyard’. Translation may involve greater levels of literality or departures from the literal meaning in the attempt to capture the authorial intentions. No surprise then if Grey’s words

‘Some pious drops the closing eye requires’

were translated by a translator as meaning that the closing eye requires some pious verses inscribed on the grave. Translations may depart from literal meaning in the attempt to capture the authorial intentions. In the same way, indirect reporting may depart from literal translations in the attempt to provide the authorial intentions. The most drastic departures from literal meanings are seen in those cases when the indirect report is confronted with an utterance whose words were not used as having literal meaning, as in the case of metaphors. To construct indirect reports based on literal meanings in these cases would end up providing unfaithful and infelicitous indirect reports. To use a term by Jaszczolt (2005), the reporter aims at ‘pragmatic equivalence’, not at semantic equivalence. This is perhaps implicit in the treatment of indirect reporting by Donald Davidson.

2. Indirect reports and language games

In this section, I will mainly discuss ideas pertaining to language games as filtered through Dascal et al.’s (1996) discussion of Wittgenstein’s language games. I quite agree with Dascal et al. that it is not easy to provide a definition of language game, and that it is best to provide examples of language games so that the notion of language game can be illustrated through exemplification. Although Wittgenstein includes indirect reports in his list of language games, to my knowledge this type of exemplification has not been properly investigated in the literature, apart from the discussion in Capone (2010a). However, Dascal et al. through their paper in which they present possibly dissonant voices without wanting to harmonize them, illustrate a kind of language game which can be called ‘presenting multiple voices’. This is similar to the language game ‘Indirect reporting’ since many of the norms governing indirect reporting make reference to the language game ‘presenting multiple voices’. Is it impossible that a language game is embedded in another language game? If it is not (as I believe), then the language game ‘presenting multiple voices’ should be embedded in the language game indirect reporting. Dascal et al. by discussing Wittgenstein’s notion of language game stress the shift (in Wittgenstein’s thought) from phenomenalism to physicalism, in the sense that language is no longer conceived of as a means of representing reality but as a means of creating (social)

reality. I quite agree with the authors on the importance of emphasizing this shift. But what consequences does this shift (or the grasp of this shift) have on the understanding of indirect reports as language games? We said that language games want to create social reality, but indirect reports seem to be anchored to language representative power. Language in such cases is used to represent occasions of language use (or events which can be called ‘utterances’). So what is the point of considering indirect reports language games if we are reluctant to stress the connection with action and favor stressing the connection with representation? My idea that a language game such as indirect reporting embeds another language game may suffice to rescue the language game from the attack we have leveled to the concept – after all, if the purpose of the language game is to make sense of the transformations of reality on the part of the reporter (and to separate one voice from another), then it is reasonable to make use of this new notion in order to account for the fact that reporting is a sort of **action** in that it transforms events in the light of the needs of hearers in the context of the reporting event. However, we do not want to confine ourselves to a derivative justification of the use of the notion of language game. If we want to see how the language game fits a conception of language in which language is used to act, we must consider narrations (as indirect reports are micro-narrations, after all) as actions. Furthermore, if we follow Tannen (1989), the language game of indirect reporting is aimed at ‘constructing’ social reality. Like actions, indirect reports can transform reality. Like actions, they can have a number of consequences. An indirect report may be part of an argumentative structure, whose aim is to justify a certain kind of action or deliberation. Thus an indirect report (of someone’s words) can be seen as a spur to act, to deliberate, etc. Seen in this light, an indirect report can become a ‘form of life’ (Wittgenstein 1953; Carapezza Forthcoming).

Another feature of language games, according to Dascal et al. is the fact that they are cooperative games. They cannot be played unless the actors cooperate in an action in which they play different roles. In what ways can this feature help us grasp the particularities of the language game ‘indirect reporting’? I propose that we conceive of an indirect report as a game involving at least three actors: the original speaker, the reporting speaker and the hearer. The hearer plays an active role in indirect reports because, many times, transformations of NPs are effected in order to favor understanding on the part of the hearer. Suppose the original speech act was

about John Campbell, suppose the reporter knows John Campbell under the modes of presentation ‘John Campbell’ and ‘The owner of the bar round the corner’ and suppose that the hearer knows John Campbell only under the guise ‘The owner of the bar round the corner’. In such circumstances, it is obvious that the reporting speaker must transform the original NP (John Campbell) into the NP with whom the hearer is familiar (The owner of the bar round the corner). The transition from one NP to another (pragmatically equivalent) is dictated by the communicative function of the indirect report. (See also Devitt 1996 for a similar argument applied to belief reports). What would be the point of issuing a report which, though faithful to the original speaker’s voice, could not be grasped and, hence, utilized by the hearer? The reporter is compelled, in these circumstances, to adapt to the hearer. The NP ‘The owner of the bar round the corner’ will express the perspective of the hearer. We could then say that the hearer plays a role in the language game – it modifies or orients the language game, has an effect over the choice of words. The recipient, in so far as she figures in the language game through a choice of words which must be regarded as potentially made by her, is one of the participants who **cooperates** in the language game. She does not speak, but her mere presence serves to perspective the game. She is present as a voice and therefore it is as if she spoke in the language game. The reporter in a sense partially, ventriloquizes the hearer.

3. Capone 2010a and indirect reports as language games

In Capone (2010a) I advanced a number of ideas on how to capture constraints on replacements of coreferential NPs in the context of direct reporting (and, in particular, in the complement that-clause). The explanation may be parallel, but not identical with the one I gave on the issue of belief reports in Capone (2008a). Such an explanation rests on the idea that replacements of co-referential NPs should not alter the speech act which the indirect report aims to report (or describe) and that the original speaker would like to see herself reported in such a way that it does not attribute her offenses, impoliteness, rudeness, obscenity, and also slurring. In other words, reporting must be done in a way that the voice of the reporter is separated from the voice of the reported speaker or, if this separation is not possible, in such a way that the original speaker’s voice is prevalent. Why should the reported speaker’s and NOT the reporting speaker’s voice be prevalent? I assume that it is a matter of

relevance. Since we are dealing with the verb ‘say’, we are happy to primarily express the original speaker’s voice and then the reporting speaker’s voice, but only if this is possible. I now succinctly sum up the main points of Capone (2010a).

The practice of indirect reports rests on the following principles:

Paraphrasis Principle²

The that-clause embedded in the verb ‘say’ is a paraphrasis of what Y said that meets the following constraint: should Y hear what X said he (Y) had said, he would not take issue with it, but would approve of it as a **fair paraphrasis** of her original utterance.

Paraphrasis/Form Principle

The that-clause embedded in the verb ‘say’ is a paraphrasis of what Y said, and meets the following constraints: should Y hear what X said he (Y) had said, he would not take issue with it, as to content, but would approve of it as a fair paraphrasis of his original utterance. Furthermore, he would not object to vocalizing the assertion made out of the words following the complementizer ‘that’ on account of its form/style.

In my paper I also discussed possible objections to the Paraphrasis/Form principle. Since this discussion will be amplified in the present chapter, I present some of the original discussion in this section.

Depending on the context, I needn’t be beholden to the original speaker’s ‘approval’ of my paraphrasis as fair, nor need I avoid manners of speech which the

² This position is somewhat reminiscent of Seymour’s (1994) treatment of in direct reports, in which reference to a translation of the reported sentence is explicitly incorporated in the semantics of in direct reports.

original speaker would shy away from. In such contexts, if John said of a person *x* that he will be coming to the party, my report to that effect is true whether I refer to person *x* politely, as John would approve of, or impolitely, as (let us imagine) my hearer would approve of. John may, upon hearing my report, demur: “Well, I don’t know why you’d call *x* a jerk but, yes, I did say he was coming to the party”. The Paraphrasis Principle and the author’s other remarks are intended to rule out contexts of indirect reporting that seem to allow this type of license with the original speaker’s words.

As I said in Capone (2010a), I am quite open to the possibility that in suitable contexts one should be able to replace an NP with a coreferential expression in the that-clause of an indirect report. However, I stick to the proposal that, in the absence of contextual clues and cues allowing us to separate the original speaker’s voice from that of the reporter, the default interpretation of the utterance conforms to the paraphrasis rules stated above. In a later section I address the issue of clues and cues as originally contributed by Dascal and I will show how it bears on the issue of the language game of indirect reporting.

4. Slurs and taboo words

An interesting phenomenon was noted by Lepore and Anderson (Forthcoming). When we report or quote uses of slurs or taboo words, the offense is assigned to the reporter (or the person quoting) rather than to the original speaker. The words that count as slurs are words which offend vast categories of people, such as black people, homosexuals, Jews, etc. So, no matter how you may want to distance yourself from the use of the word ‘negro’ in reporting such a use you are also liable to be blamed for the use. Lepore and Anderson are adamant that there is something offensive in reporting a slurring word. This is not to say that there are no contexts where the offense is mitigated or nullified (e.g. the academic context in which one discusses a phenomenon scientifically and goes to great pains to show that one deplors such uses). However, in ordinary contexts reporting slurs or taboo words is not convenient. Lepore and Anderson reject the idea that slurs be considered presuppositions or be explained by resorting to the notion of conventional implicature. Verbs of saying such as those involved in indirect reporting, in fact, are

plugins to presuppositions and thus we should predict that the presuppositions are filtered out by such contexts. The notion of conventional implicature, being semantic, cannot explain why slurring persists despite embedding in verbs of saying in negative structures. Conventional implicatures, in fact, appear to be negated:

(4) John did not say that Mary is French but brave.

It appears that neither John nor the reporter of (4) is committed to there being a contrast between being French and being brave (but it is presupposed that at least someone might hold this contrast). However, slurs and taboo words persist despite embedding in verbs of saying.

I accept that reports of slurs or taboo words are often embarrassing, though I do not completely agree with Lepore and Anderson that only the reporter is perceived as being responsible for the slur. I am persuaded that mainly the reported speaker is perceived as responsible for the slur and that at most responsibility is shared between the reporter and the reported speaker. I would also like to insist that even if the addition of prefatory speech aiming at providing justifications for the report and at softening the hearer's judgments cannot completely nullify the offence, it can at least serve to mitigate it.

Lepore and Anderson insist that neither conventional implicature nor presupposition can explain the embarrassment caused by reporting slurring words and finally opt for the explanation according to which there is a prohibition against using slurs which applies both in the case of the original speaker and of the reporter. I applaud this explanation, as this is an explanation which considers the practice of reporting a speech act a language game subject to rules of use. It is not semantics (either in the form of conventional implicatures or of presuppositions) which prevents us from reporting slurs easily, but it is a rule of use. We must know what uses are licit or not and thus, to master the practice of the language game 'reporting speech', we must know what uses are licit or prohibited. The rules and uses in question are not semantic, but are **societal**. We must know what practices society allows and what practices it bans. Thus using (or not using a word) is ultimately a matter of knowing societal, rather than linguistic uses.

If the considerations by Lepore and Anderson were completely accepted in their current form, this would have severe implications concerning my stance to the

principles involved in the language game ‘reporting speech’. In fact, in my 2010 paper on indirect reports (Capone 2010a), I claimed that whenever strong contextual clues allowing us to distinguish voices within the same speech act are missing, considerations of processing efforts militate against having two voices in the same report. The interpretation that is most relevant is the one that attributes the voice expressed in the *that*-clause of the report to the original speaker, since ultimately the aim of reporting is to say what an original speaker said. Having two voices involves greater processing efforts which are justified if the ultimate effect of the utterance is to let the hearer know how to separate the reporter’s voice from the reported speaker’s one – that is to say if numerous contextual clues are present and allow separation of voices.

So if Lepore and Anderson are correct that the slurring words are attributed to the reporter, rather than to the reported speaker, my principles would have to be abandoned. However, I opted for a weaker version of Lepore and Anderson’s theory, according to which both the reported speaker and the reporting speaker are perceived as slurring. My own theory actually predicts that the greater offence is attributed to the reported speaker.

5. Default interpretations and modularity of mind.

In Capone (2010b), I explored the idea that explicatures are the result of modular processes. In particular, I argued that the view by Kasher (1991) according to which only the pragmatic processes involved in understanding speech acts are modular needs to be reassessed. Kasher believes that cancelability (which is one of the characteristics of conversational implicatures) is an obstacle to seeing the interpretative processes involved in pragmatic inferences as modular. In fact, modular processes must be both mandatory and encapsulated. However, the fact that inferences can be aborted when we have access to a body of knowledge, shows that implicatures are not mandatory and, furthermore, require interaction with vast archives; hence they cannot be encapsulated. In Capone (2009), however, I argued that explicatures are not cancellable and this seems to be in conflict with Kasher’s (1991) ideas. In Capone (2010c), furthermore, I argued that the processes involved in calculating explicatures are encapsulated (often requiring ‘modules on the fly’ to search information). In Capone (2010d) I argued that pragmatic processes involved

in conversational implicatures can be of two types – modular or non-modular. Following ideas by Cummings (2009) I accepted that at least some cases of conversational implicatures involve the interplay of vast archives of information. However, modular pragmatic processes usually provide the propositional forms which are **then** utilized in non-modular processes. Now, this means that default inferences usually arise due to modular processes – these can become the final messages or otherwise be aborted or integrated by information accessed through vast archives (encyclopedic knowledge). The integration of propositional forms obtained by default inferences through access to vast archives is itself constrained by the Principle of Relevance. Thus, we will require that contextual effects and processing efforts are to be kept in balance while the integration of default propositional forms takes place.

Now, returning to the issue of indirect reports, what are the consequences of this modular approach? One of the consequences is that we consider the default interpretation of an indirect report one in which the voice of the original speaker (the reported speaker) is presented UNLESS there are ways of distinguishing voices (there can be at least two voices, but even more than two voices if we consider that reports may come as the result of chained indirect reports) and of assigning voices to the respective participants. In Capone (2010a) I claimed that processing efforts are involved in selecting the voice of the original reporter. Surely, it must be admitted that a report in which there are two inseparable voices is one that requires greater processing efforts (as one cannot easily settle the question ‘Whose voice is this or that one?’). I admit that this way of resolving the problem is partial and not definitive. One must not only show that processing efforts play a role in deciding whether one or rather two voices, are present in the micro-narration of the indirect report, but ideally one should be able to demonstrate that we have a pragmatic explanation of why the original speaker, rather than the reporting speaker, is selected as the voice that counts. And here the problem must be framed in terms of contextual effects, rather than in terms of processing efforts, since if the original speaker’s voice prevails, then we have additional information on the perspective of the original speaker – and now we remind readers that indirect reports are micro-narrations about the original speakers, about events in which the original speaker and not the reporting speakers were involved. Thus, if the modes of presentation of NPs used in indirect reports are those used by the original speakers, we have additional

information on their point of view, on their language, of their mental processes, and also on the context of utterance. These make contextual effects larger, following Sperber and Wilson (1986).

6. Dascal and Weizman (1987) on clues and cues.

Dascal and Weizman (1987, republished in Dascal 2003) is a superb exposition of the issue of clues and cues, which I want to put to use in the understanding of the logic of indirect reporting. Dascal and Weizman, following a tradition going back to Searle (1979), notice that understanding a speech act is often a matter of filling the gaps left there by the speaker by using pieces of information available in the context (whether intended as the specific situation of utterance or background and cultural information having a bearing on the utterance). Since texts may often depart from literal meanings in substantial ways, requiring not only filling of deictic elements but also drastic revision of the literal meaning, Dascal and Weizman are ready to admit that two types of instructions are needed to make sense of texts. These are called contextual 'clues' and 'cues'. These are differentiated in a functional way. Cues allow readers or hearers to notice drastic departures from literality (e.g. Is a lexeme used metaphorically? Is the speaker conveying irony? etc.) allowing contextual clues to provide specific solutions to the general problem addressed by cues. So, the questions introduced by cues are normally yes-no questions. The questions introduced by clues are more open, involving a search for items capable of filling lacunae in a text. As Dascal and Weizman cogently say, a cue problem soon turns into a clue problem. This interplay of cues and clues is at the basis of understanding a text. A practical method invented by Dascal and Weizman for distinguishing cues is to ask informants to transform a text. Those elements which are radically transformed in paraphrasis signal a cue problem. This is clearly relevant to understanding the logic of indirect reports, since paraphrasis is involved in reports and indirect reporting is perceived as deviant in some ways if the utterance is reported literally without taking into account the cues and clues which lead to meaning augmentations and legitimate departures from literal meanings. Indirect reporting is ultimately a way of checking whether the interplay of cues and clues has led to plausible meaning augmentations, because if it has not, then the report is not legitimate. I propose that indirect reporting is closely connected to the issue of cues

and clues as presented by Dascal and Weizman. There is another point of intersection between Dascal and Weizman's considerations on cues and clues and the logic of indirect reports. Dascal and Weizman discuss in detail various types of clues and distinguishes between:

Clues related to extra-linguistic specific context;

Clues related to meta-linguistic specific context;

Clues relating to extra-linguistic shallow context: general assumptions about the features of a given set of situations;

Clues relating to extra-linguistic background knowledge: general knowledge about the world;

Clues relating to meta-linguistic background knowledge: general knowledge about the functioning of verbal communication.

Clues related to meta-linguistic specific context play an important role in indirect reports. We have, in fact, often said that understanding a report fully is a matter of separating the original speaker's voice from that of the reporter. Thus a clue that tells you that a certain word is part of the linguistic repertoire of a certain person (suppose that word is idiosyncratic to that person) will also allow you to separate the original speaker's voice from the reporter's voice. But first of all, the idiosyncratic word may constitute a cue allowing you to notice that there is an interpretation problem relating to indirectness. Then the cue problem will turn into a clue problem and the cue/clues will allow you to sift the original speaker's voice from the reporting speaker's voice. But now the question arises whether the interpretation process pertaining to separating voices in indirect reports can be included in the more general rubric 'noticing a discrepancy between what the speaker literally says and what the speaker's meaning is'. My answer is positive. Although the cue does not in this case allow you to detect a drastic departure from literal to intended meaning, it will allow you to establish a more accurate structure in the report and to fill the lacunae thanks to contextual clues. So, in a sense, the interpretation problem posed by indirect reports is a sub-case of the more general case discussed by Dascal and Weizman.

There are further parallels between the discussion in Dascal and Weizman (1987) and the case of indirect reports. Consider the example discussed on:

father of fathers of.

This expression is used to express the concept: the original cause of. According to Dascal and Weizman, the departure from literal meaning must be detected through a cue of indirectness:

A cue for indirectness is to be found, if the reader employs her meta-linguistic shallow knowledge and, via the specific meta-linguistic acquaintance with Biblical style (to which the literal Hebrew expression belongs), notices an unexpected register shift. (Dascal and Weizman 1987, republished in Dascal 2003, 189). (See also Weizman and Dascal 1991 on extralinguistic shallow knowledge enhancing associations with the notion of ‘fighting family’).

Now, following these ideas by Dascal and Weizman, I propose that meta-linguistic specific context provides cues and clues allowing hearers to separate voices in an indirect report. There are many ways in which a reporter can allow hearers to recognize voices: they can imitate the voice quality, they may use items of vocabulary which are idiosyncratic to a certain speaker (including the reporter), they can use stylistic features that are recognizable as belonging to a certain well-known author, etc. (See also Recanati 2001). We should add that interpretative problems increase in complexity if the reported speaker in turn embeds somebody else’s voice in his own voice (see cases of mixed quotation).

The ideas by Dascal and Weizman were taken up by Hirsch (2011), who applied them to humor. Typical cues for humor are discussed by Hirsch: script opposition (the violation of expectations), framing (jokes appear to have repeatable structure, usually a single scene terminating with a punch-line), word play and nonsense. For the sake of space I cannot go into this, but needless to say these ideas are very useful when it comes to identifying discrepancies between literal and intended meaning in indirect reports. Some parts of an indirect report can be humorous and it may be important in such cases to distinguish between the voice of the original speaker and that of the reporter. Who is being humorous? Although, I will not specifically discuss humor and indirect reports, I mention this possibility as part of the general task we

are confronted with of separating and specifying the voices expressed in an indirect report. Obviously, we need cues and clues to separate such voices.

6.1 Applications of cues and clues.

In this section I will apply the notion of cues and clues to three important cases of indirect reports. Needless to say, I will keep the discussion short for the sake of time, but I imagine that a number of other cases need to be discussed or taken into account. A case that is of great theoretical importance is an indirect report with implicit **translation**. Sure, by now we have arrived at the plausible tenet that paraphrasis is involved in indirect reports. Paraphrasis may involve shortening (summing up) or even expanding the report. What makes an indirect report legitimate is the extent to which we are ready to re-express the original voice without distortions of the message or of the form of the message. In case of reports with (implicit) translation it is implicit that the paraphrasis was reached through a translation (the question ‘whose translation?’ is not to be easily dismissed). Consider the report:

(5) Putin said that any American attempt to increase the nuclear arsenal will be considered as a threat to the talks on disarmament.

Now it is clear enough that (5) is a paraphrasis of what Putin said – and this may well consist of an abridgment and of a translation. The translation is obviously one from Russian into English. So we may well accept that Putin’s words were very different from the ones used in the indirect report. It is true that a polyphonic reporter may well utter the sentence (5) by using a recognizable Russian accent and may even try to imitate the specific quality of the voice of the Russian leader. However, there may be strong cues telling us that there may be a divergence between the words used by the reporter and those used by Putin and clues leading the hearer to guess that Putin spoke Russian when he uttered the message (Suppose Putin prefers to speak Russian rather than English, which he may know well, due to patriotic reasons). In this case, the following types of clues may be relevant:

Clues relating to extra-linguistic shallow context: general assumptions about the features of a given set of situations;

Clues relating to extra-linguistic background knowledge: general knowledge about the world;

In particular, the clues guiding our interpretation are knowledge of the general fact that Russians normally speak Russian and that the Russian political leader due to patriotism may want to speak Russian when addressing foreign policy (perhaps even before foreign journalists).

The second type of example pertains to non-literal uses. Consider ironic uses. John and Mary are in the library, studying. It is almost time to leave the library and go to dinner. John says: ‘Are you staying here?’ and Mary replies:

(6) Yes, I will stay here all night long.

Now, we want to ask whether Fred, who was near them and overheard the conversation could legitimately report:

(7) Mary said that she will stay in the library all night long.

In a sense, it may not be legitimate to report Mary’s speaker meaning by a report of a literal utterance, if what is required of the reporter is Mary speaker’s meaning. So, whether or not an indirect report of an utterance is legitimate if it only reports the literal meaning of the sentence uttered very much depends on the requirements of the context. If the report was elicited by someone in need of the speaker’s meaning, it would be illegitimate to report the literal meaning. In general, there should be

constraints preventing reporters from only reporting literal meanings of sentences. These constraints come from general Gricean principles or Relevance Theory considerations (the Communicative principle of Relevance): one should avoid ambiguities and NOT put the hearer to undue and unnecessary processing efforts. They also specifically come from accepting the Paraphrasis Principle which I discussed in Capone (2010a) and reported here in a previous section. One should notice that there is a tension between the Paraphrasis/Form Principle and my claim that it is not standardly legitimate to report literal meanings, when ironic messages are at stake. So how can we resolve this tension? The tension amounts to this: a level of literality is needed to prevent distortions and attributions of slurs, foul language, sexist language, taboo words, offensive language in general to the original speaker, when, in fact, these should be attributed to the reporter; a level of literality cannot be tolerated when the speaker's meaning diverges from the literal meaning in a drastic way as in ironies. Yet, the tension disintegrates if we consider that both in cases of NPs and of whole utterances speaker meaning is involved. No one prevents me from using the NP the original speaker used in his speech by using it both literally and as being speaker-intended. In other words, in some cases literal meanings are also speaker-intended. Thus, whether we report what the speaker said by using parts of what he literally said (and speaker-meant) or we have to drastically alter what he said to capture the speaker's meaning, we ultimately report the speaker's meaning.

Another issue I want to tackle in connection to reporting of literal meanings is whether one really cannot find ways to report literal meanings and do so in a way that is considered acceptable at least in some circumstances. We have said that the reason why we intend to interpret indirect reports as reports of speaker's meaning is that doing otherwise generates ambiguities that cannot be easily resolved. And one is under the constraint to avoid ambiguities and to put the hearer to as little processing effort as possible. However, if processing efforts are balanced by rich cognitive effects, then it may be acceptable to report literal meanings. So even if we admit that this is not the general practice of reporting, in some cases where knowing what the speaker literally meant is of importance to the hearer, the consideration of processing costs is put aside in view of the richer cognitive effects. Suppose that something important and crucial hinges on what Mary literally said and this is made clear in the reporting context. Suppose further that there are rich cues and clues allowing us to assess that the reported statement is a literal and faithful reproduction

of the original speaker's words. Then we have the opposite process of what was described by Dascal and Weizman. In the original context cues and clues allow us to detect an interpretation problem and to construct the speaker's meaning. In the reporting context, there are rich cues and clues allowing us to detect an interpretation problem (in particular that the general practice of interpreting indirect reports is not applicable) and assign the words of the report a literal meaning status. What is clear is that in some contexts, this is possible and it is possible due to the existence of rich cues and clues.

The last case I want to discuss, in connection with cues and clues is an example proposed by Tannen (1989). This example is interesting because it corroborates what we have said so far, that is that reported speech is often a transformation of the original words, which requires cues and clues for interpretation. Tannen like me proposes that intended meanings, considered as 'constructed', should be at stake in indirect reports, and that interpretative problems arise when the recipient fails to reconstruct an utterance's intended meaning.

The example is the following. Two sisters talk on the phone (let us call them A and B). A reports what their mother said about B: she criticized B for not returning home. Instead, she apparently did not criticize A for not returning home from college. B notices the illogicality of this situation, as A who is a college should be under a greater obligation to return home during holidays. B resents her mother's criticism, and apparently seems to take what her mother says literally. However, Tannen proposes that the mother for tact reasons avoids criticizing A directly but conveys to her her disappointment by criticizing B, expecting A to reason that if her mother has reasons for criticizing B, then she also has reasons for criticizing A. Apparently, neither A nor B grasp their mother's intended speech act and, thus, in reporting it, A concentrates on the literal meaning. Yet, there are cues and clues sufficient for signaling an interpretative problem and for solving it by assigning a specific interpretation of the speaker's intentions. The question which we may address now is whether these cues and clues are preserved by the telephone conversation between the two sisters. We have got situation S, where there is a telephone conversation between the mother and daughter A and situation S1, in which there is a telephone conversation between sister A and B. It is certainly possible, that certain cues such as quality of voice were lost when the first conversation was reported to B. So it is

possible that when the mother told A that it was ok for her if she did not return home, her voice was colored in a certain way and such quality could constitute a cue. It is possible however, that this cue was lost in the reporting of the conversation. However, the cue and cues constituted by the criticism of the other sister for a conduct which was similar to the conduct by the other sister who was not actually criticized do not disappear. The problem, however, is that sister B is not able to perceive such a cue, is not able to reason on the discrepancy between the mother's reasons for criticizing her conduct and her reasons for approving a similar conduct by her sister. Basic rationality principles could be easily conducive to a reasoning about what the mother's intended meaning is. So why is it that the two sisters were not aware of the intended meaning? Why is it that the indirect report only takes into account literal meanings? Tannen says that in American English indirect reports are taken as reporting literal meanings. But this is to be excluded by what I argued before in this same section. Tannen is ambivalent on this; on the one hand, she takes American speakers to orient to literal meanings in indirect reports; on the other hand she says that

I am claiming that when a speaker represents an utterance as the words of another what results is by no means describable as "reported speech". Rather it is constructed dialogue. And the construction of the dialogue represents an active, creative, transforming move which expresses the relationship not between the quoted party and the topic of talk but rather the quoting party and the audience to whom the quotation is delivered. (Tannen 1989, 11).

But we have already seen that the picture of indirect reporting which does not consider transformations, voicing, cues and clues is deeply flawed. Thus, I take Tannen to express reservations for a notion of indirect reporting which is close to verbatim reports, but not to the notion of indirect reporting which we have defended in this chapter. The idea that **construction** is involved is also a familiar one for us, since we have already accepted that explicatures must be part of indirect reports, that an implicit reference to translation is sometimes made, that the reported words can be summed up or even articulated in a more precise and elaborate manner. Construction

also involves constructing a framework for separating voices. It also finally involves assigning an illocutionary intention to a speech act. Cues and clues, in the sense of Dascal and Weizman (1987), feature prominently in this picture of how indirect reports are constructed and deconstructed. The construction work involves taking into account a number of contextual elements which first of all tell you how something has to be taken and then allow you to assign specific content or voices to the indirect report.

Conclusion

In this chapter I have argued that indirect reports often involve transformations of the message which was uttered in the reported utterance. There is a limit to the quality and number of transformations applicable, and such a limit can be found in norms that regulate the language game 'indirect reporting'. I argued that indirect reporting is a societal practice, which involves societal pragmatics considerations along the lines of Dascal and Weizman (1987) and Mey (2001). Societal pragmatics must be allied with cognitive pragmatics, since the norms implicit in the practice of indirect reporting can often be deduced by cognitive principles like, for instance, the Cognitive Principle of Relevance. Since a theory of how indirect reporting works in conversation is also a theory of communicative practice, it follows that the Communicative principle of Relevance is also at work. This paper leaves some matters unsettled, though. It would be useful to consider the interaction between indirect reports and the theory of quotation. I leave this matter for another day.

Chapter 9

Knowing how and pragmatic intrusion.

Abstract.

In this chapter I try to apply pragmatic intrusion to the controversy of whether knowing how can be reduced to knowing that. I argue that pragmatic intrusion can resolve a number of problems encountered in the work by Stanley and Williamson (2001).

Introduction

I start this chapter with some programmatic statements. We have seen how Grice's pragmatics was assimilated and put to use by linguistics in the course of this and the previous century. The assimilation of Grice's ideas almost amounted to a revolution and spurred an enormous amount of work on this topic. However, now the time has come for a new period of research, in which Grice's philosophical ideas are applied (back) to philosophical topics. We have sufficient evidence showing that this can be done successfully (see for example the brilliant paper by Douven 2009). The quantity of work on propositional attitudes and pragmatics indicates that this type of program can be successfully executed and should, therefore, be executed. It would not be a bad idea to extend the topics to

which we can apply our ideas about pragmatic intrusion and explicatures. I am persuaded that knowing how is one such topic in which pragmatics can be applied to philosophy in the attempt to resolve a philosophical problem. However, since ‘knowing how’ attribution is also a linguistic practice, we should also find some harmony between philosophical and linguistic descriptions. It is, therefore, natural to want to explain a number of inferences triggered by knowing how utterances in terms of pragmatics.

In this chapter I shall deal with the pragmatics of knowing how utterances. Since ‘knowing how’ vs. ‘knowing that’ has received detailed treatment in the philosophical literature (see Fantl 2009 for a good review), we cannot but pay attention to what philosophers have said. However, in the main, I want to stress what is needed for a theory of communication, rather than a theory of knowledge. Since a theory of knowledge and a theory of communication intersect at various points, this is not easy to do. For example, one reason why we utter sentences such as “John knows that p” is to inform the hearer that p. Hearing a sentence like ‘John knows that p’ the hearer is entitled to infer that p, provided that he trusts me and he trusts John. Nevertheless, I assume that epistemology and a theory of communication are different projects, with different aims, and it does no harm to reveal my bias towards a theory of communication. In starting this chapter, my hope is that we shall throw light on pragmatic inferential processes involved in understanding apparently simple knowing how utterances. Essentially, I want to reflect on what is being communicated and on what inferential processes are triggered when knowing how utterances are proffered. I will claim that such inferential processes are responsible for what is being communicated at a truth-conditional level.

Indirectly, my inferential approach serves to throw light on many problems that bear on the distinction between knowing how and knowing that and on the possibility of reducing knowing how to knowing that. My inferential approach shows that, despite recalcitrant data leading to the opposite view, the right interpretation of these data will, in fact, lead to the idea by Stanley and Williamson that knowing how is a species of knowing that.

I start the chapter with the exposition of Stanley and Williamson's (2001) paper according to which 'X knows how to p' must be analysed as 'X knows that there is a way w and that he can p in this way'. I then move on to the discussion of opacity, which has been claimed to constitute another reason for the asymmetry between knowing how and knowing that. I argue that opacity characterizes both knowing how and knowing that and answer some objections. I also discuss an objection to Stanley and Williamson's treatment of Gettier's problem. My analysis is that knowing how can be Gettiered. I then consider objections by Sgaravatti and Zardini to an assimilation of knowing how to knowing that based on negativity arguments and closure principles. I argue that these objections are not well-founded and can be dissolved through pragmatic intrusion. Subsequently, I discuss knowing how and modularity of the mind, and claim that modular considerations need not be inimical to a unified treatment of knowing how à la Stanley and Williamson. I argue that the considerations on contextualism and knowledge by Keith De Rose are also applicable to knowing how, in addition to knowing that. I also argue that Igor Douven's considerations on the pragmatics of belief apply to knowing how as well as to knowing that.

In Part 2, I present a brief section on the notion of pragmatic intrusion, which will be of help in understanding the issue of knowing how. I provide some of the standard examples of pragmatic intrusion. Then, I deal with a number of uses of the 'knowing how' construction, and I focus, among other things, on the asymmetry between knowing that and knowing how, noticed by Snowdon. While one can deduce 'I know that p' from 'X knows that p', it is not clear that one can deduce 'I know how to p' from 'X knows how to p'. In section 2.2 I consider a famous objection by Rumfitt to Stanley and Williamson's treatment and claim that it can be dissolved through pragmatic intrusion. In section 2.3, I consider various inferences in context and also argue in favor of pragmatic intrusion, considering a weak form of contextualism. I also consider interpretative problems with quantifiers and implicit arguments. In section 2.4, I argue in favor of a unified treatment of knowing how and knowing that, through pragmatic intrusion. I consider two alternative hypotheses.

1. Stanley and Williamson on knowing how.

Stanley and Williamson (2001) question Ryle's (1949) considerations on the distinction between knowing how and knowing that, a distinction which has engaged many philosophers of mind and of language. In addition to this, they question Ryle's claim that an action of any kind requires a previous act of thinking, to count as intelligent. In this chapter, I will not dwell on this second, and important issue, but I will concentrate on whether or not knowing how can be subsumed into knowing that. Stanley and Williamson question Ryle's distinction on the basis of linguistic arguments – a move which has caused Michael Devitt's perplexity (since he claims that metaphysic claims should be settled by metaphysic arguments (Devitt Forthcoming)). Yet, it is not true that Stanley and Williamson simply deal with the linguistics of 'knowing how'; they also deal with various connected phenomena which are typical of propositional attitudes (e.g. Gettier's problem; the opacity problem). It simply turns out that many of the problems encountered in the discussion of the distinction between knowing how and knowing that can be resolved – if not dissolved – through the recognition of pragmatic intrusion. In this chapter I want to claim that one cannot understand the issue of knowing how independently of the understanding of the phenomenon of pragmatic intrusion.

So I am sympathetic in principle to Stanley and Williamson's (2001) 'linguistic' treatment of knowing how, since linguistics and philosophy in this case intersect to a greater extent than in a number of different issues.

To make complex things simple, Stanley and Williamson take cases of 'knowing how' to be cases of indirect questions (embedded questions) (In a later version of this paper, in press, Stanley cites Higginbotham (1993)). In this perspective, a sentence such as 'Ann knows how to ride the bike' has the same truth-conditions as 'Ann knows that there is a way *w*, which is a way for Ann to ride the bike'. To be more explicit, this must be interpreted as 'Ann knows that there is a way *w* such that she can ride the bike in way *w*. This proposal admits three kinds of pragmatic intrusion: a) assigning the 'some rather than all' interpretation to the knowing how sentence (Ann knows some way (not all ways) which is a way for her to ride the bike); assigning a deontic modal interpretation to the infinitival construction (Ann knows how to drive the bike → Ann knows

there is a way w such that she can ride the bike in way w); assignment of a ‘de se’ interpretation to PRO (Ann knows that w is a way for her to ride the bike; NOT: Ann knows that w is a way for one to ride the bike; I will return to this point later on).

The proposal becomes more complicated when it comes to modes of presentation. In line with considerations about propositional attitudes, given that ‘know’ is like ‘belief’ in giving rise to opacity problems, one had better introduce the notion of modes of presentation of the reference. The reference can present itself to the knower in different ways. To account for the cases in which knowing how is equivalent to having a skill, Stanley and Williamson introduce the notion of a practical mode of presentation. (Basically in addition to being told how to p , you can be shown how to p .)

1.1 Opacity.

Those who want to emphasize the contrast between knowing how and knowing that typically invoke two characteristics of knowledge claims: opacity and Gettier’s problem (Gettier 1963). Knowledge claims are typically opaque; in other words, substitutability of identicals (Leibniz’s law) does not work inside the scope of a knowledge/belief operator. Suppose that Elisabeth = the Queen of England. Then in a sentence such as (1)

(1) Elisabeth greeted the Polish ambassador

one can replace ‘Elisabeth’ with ‘The Queen of England’ without changing the truth conditions of the sentence. However, in a sentence such as (2)

(2) Mary knows that Elisabeth greeted the Polish ambassador

one cannot freely replace ‘Elisabeth’ with ‘The Queen of England’ (in case Elisabeth = the Queen of England) UNLESS one also knows that Mary knows that Elisabeth = the Queen of England. After Salmon (1990), one tends to

interpret opacity pragmatically. Truth-conditionally, the sentence ‘Mary knows that Elisabeth greeted the Polish ambassador’ is the same as the sentence ‘Mary knows that the Queen of England greeted the Polish ambassador’, however they differ in the implicated import. Someone who utters (2) somehow commits herself to accepting that Mary knows Elisabeth through the mode of presentation ‘Elisabeth’ and NOT through another mode of presentation. Thus she might not accept that Mary knows Elisabeth under the mode of presentation ‘The Queen of England’ and, therefore, she might not accept the statement (3)

(3) Mary knows that the Queen of England greeted the Polish ambassador.

Those who want to say that knowledge-how and knowledge-that claims are completely different have the opportunity to deny that opacity applies to knowledge-how. Suppose that John attends regularly some Karate classes and knows how to make a move that involves three steps; let us suppose that there is also a type of dance, let us call it the Cuckoo dance, involving these three steps. Then we should describe John, who knows how to practice Karate, as someone who knows how to engage in the Cuckoo dance. Yet, he would hardly be willing to describe himself as someone who can practice the Cuckoo dance. He would answer negatively the question whether he can practice the Cuckoo dance. But why should it matter whether John is willing to accept that he can practice the cuckoo dance? We know that, if we bothered to explain to him what the cuckoo dance is like, he would be able to do it.

There does, therefore, ‘prima facie’ appear to be an asymmetry between knowing that and knowing how. When we come to knowing that, the subject’s judgements about what he knows are important. After all, usually one knows that he knows that p (but see Williamson 2002 for the claim that one who knows p need not know that he knows p). If the subject’s judgements are important, then one hardly knows that p UNLESS one knows that he knows that p and, if p is presented under a different mode of presentation, the risk is there that the subject will not accept that he knows that p under mode of presentation q, because he has introspective access to his knowledge state and he does not find there the proposition q.

Now, the contrast between knowing how and knowing that, in terms of opacity seems to depend on two facts: luminosity; and pragmatic intrusion. Concerning luminosity, I am inclined to say, following Williamson (2002), that not in all cases in which one knows p one knows that one knows p , albeit normally luminosity is implied. Concerning pragmatic intrusion, we must remember that opacity was heavily dependent on that. There may be pragmatic reasons why pragmatic intrusion occurs to a greater extent in knowing that claims than in knowing how claims. Knowing how claims, in fact, are normally connected with practical problems. Why do I know whether John knows how to p ? Presumably I want to know whether John knows how to p because I need to p . And if I know that John knows how to q and that knowing how to q is equivalent to knowing how to p , then I can use John's know how for the purpose of resolving my practical problem. Knowledge that claims, on the contrary, are not connected with practical problems, but with inferential knowledge. Often I need to know that p , because p , together with q , tells me whether n and I need to know whether n is true in order to do F . (And I can know p if I know that John knows that p). Opacity is clearly a problem in inferential knowledge, since it is a limit to what can be inferred given a knowledge claim. Surely it would be of help to me to know that Mary knows that the Queen of England will be here today – I would know that there is a likelihood that the traffic will be jammed and I will avoid driving in the city centre. Yet, there is no way to extract this information from knowledge that Mary knows that her friend Elisabeth will be here today.

So my conclusion is that knowing that and knowing how behave differently with respect to opacity not because they are inherently different claims, but because they serve different purposes. It is the function of the claims that makes them appear as if they were different with respect to opacity.

1.2 Gettier's problem

Many have found it reasonable to assume that if knowing how is not completely different from knowing that, then it is vulnerable to Gettier's problem. In short, Gettier's problem consists in the fact that a proper knowledge claim must not admit devious types of justification. For instance, the following case does not

count as knowledge. John sees Billy drive a new Mercedes. He concludes that he is the owner of a new Mercedes. Thus, he is able to make the inference that either Billy or Fred bought a new Mercedes. It then turns out that Fred bought a new Mercedes, but Billy didn't (he was simply trying Fred's car). Then by pure luck, it is true that Fred bought a new Mercedes; hence it is true that either Fred or Billy bought a new Mercedes. John believed that either Fred or Billy bought a new Mercedes; he was somehow justified in coming to that belief. So should we say that he knew that Fred or Billy bought a new Mercedes? The standard answer has been negative and epistemologists have pondered on the nature of the justification. In short, the justification must be non-devious.

Now, can knowing how claims be Gettiered? Stanley and Williamson (2000) answer this question positively and contrive a case to show that knowing how can be Gettiered. Suppose that John wants to learn how to fly a plane and does that through a flight simulator. Wicked Fred changes the simulator's programs intending to disorient John, thus preventing him from learning how to fly a plane. Yet, despite his evil intentions, by mere luck the simulator activates the right programs, designed to train a learner how to fly a plane. When he finishes his simulation courses, John is actually able to fly a plane and everyone admires him for his skill. Should we say that he knows how to fly the plane? If he knows how to fly a plane, he knows that by mere luck. Surely the success is fortuitous and even if AFTER he actually proves that he is able to fly a plane, we are willing to grant him know-how, we would be reluctant to let him fly a plane if we knew what evil Fred has done to the simulator. After seeing what Fred has done to the simulator, we would predict that John will not be able to fly the plane. We would not be willing to let him fly the plane because we think he has not acquired know-how.

It has been objected (Poston 2009) that, appearances notwithstanding, John has know-how. He knows how to fly the plane, and that is all that matters to show that knowing how cannot be Gettiered. Yet, it would be reasonable to suppose that when John knows how to fly a plane, he knows all the kinds of steps involved in flying a plane and he knows of all ways w in which one can fly a plane, that w is a way for him to fly the plane. How can he know that w is a way for one to fly the plane if he learned that w is a way to fly the plane through an evil instructor? If

knowing how has to be learned, it must be learned in the proper way. The case is analogous to knowing that. Suppose that Mary tells me that Fred is in Paris. Fred happens to be in Paris. Yet Mary was evil enough to tell me a lie. Did I really know that Fred was in Paris? The justification for my belief was devious, thus I did not really know. In the same way, when an evil instructor tells me (or shows me) that w is a way to fly a plane, hoping to induce me to make a mistake and, unaware to him, w proves to be an efficient way to fly the plane, then it is by mere luck that I can fly the plane, as my justification for believing that w is a way for me to fly the plane was devious.

1.3 Sgaravatti and Zardini against the assimilation of knowing how to knowing that.

Before moving on to a deeper treatment of knowing how vs. knowing that, I would like to consider some interesting objections by Sgaravatti and Zardini (2008) to subsuming knowing how into knowing that. The objections are intelligent and worth-while discussing. The analysis according to which knowing how is subsumed into knowing that, following Stanley and Williamson, amounts to more or less the following:

John knows how to ride the bike = def John knows that there is a way w which is a way to ride the bike.

This analysis is still incomplete, in that it requires the introduction of a modal element, but this will not change the gist of our discussion.

Now, Sgaravatti and Zardini argue against Stanley and Williamson on the basis of negativity. Consider (17)

(4) John knows how to square the circle.

Translated through Stanley and Williamson's semantics, this would come to:

(5) John knows the answer to the question how he can square the circle

or

(6) John knows that w is an answer to the question how he can square the circle.

However, Sgaravatti and Zardini argue, John knows a perfect answer to the question how to square the circle and that is 'In no way!', but nevertheless he cannot square the circle. When the negativity argument is introduced, Stanley and Williamson's treatment faces a serious problem.

If I were Stanley and Williamson, I would not be worried by the negativity problem. I assume their answer would be that if one answers 'In no way!' to the question how to square the circle, then one has no answer to the question how to square the circle. If you were to ask me to make 8 loaves of bread out of just two loaves, or how one can make 8 loaves of bread out of just two loaves, I would not be able to provide a satisfactory answer to your question. Knowing how involves providing a satisfactory answer to a (practical) question, and not just an answer whatsoever. If providing an answer whatsoever to the question how I can drive a powerful motorbike is what enables me to drive a motorbike, then, I (who cannot ride a powerful motorbike) could qualify as knowing how to drive a powerful motorbike. But it is not an answer whatsoever, that is required, but an answer that allows me to carry out a task. In Sgaravatti and Zardini's case, could we honestly say that the answer 'In no way!' provides a satisfactory answer to the question how to square the circle? If it does not, then we should not be surprised if one cannot manage to square the circle, even if one is able to provide an answer (which turns out to be inadequate).

If anything, Sgaravatti and Zardini's contribution is to clarify what is or must be intended when one says that w is the answer to a question. Presumably, the notion of Relevance (à la Sperber & Wilson 2004) may be of help. Relevance

is defined (by Sperber & Wilson) in terms of contextual effects and processing efforts. There are three ways in which a newly presented piece of information may interact with the context to yield a contextual effect: (i) it may combine with contextual assumptions to yield a contextual implication (that is a logical implication derivable neither from the contextual assumptions nor from the new information alone); (ii) it might strengthen an existing assumption; (iii) it may contradict and lead to the elimination of an existing assumption.

Now, let us reconsider the question whether the answer ‘In no way!’ is relevant to the question ‘How can one square the circle?’. Does this answer create cognitive effects which justify the effort of vocalizing it? One thing the answer does, surely; that is, it questions the question and categorizes it as a question that cannot be answered. But then ‘In no way!’ is NOT an answer to the question, but a challenge to the question, amounting to, more or less, “Why did you ask me such an absurd question?”. Then, if it is not an answer to the question, it is not in conflict with the fact that nobody actually possesses this know-how. I assume Relevance Theory – and in general theories about communication – have a role to play in clarifying this not uninteresting case posed by Sgaravatti and Zardini. At this point, I want to ask the question whether the one we have just encountered is a case of **pragmatic intrusion**. Pragmatic intrusions are cases in which a pragmatic inference allows truth-evaluability. If I remember well, the cases of pragmatic intrusion discussed by Carston (1999) were exactly cases in which, unless pragmatic intrusion was posited, one ended up with a logical form which could lead to contradictory interpretations. Consider one of those examples:

(7) If the King of France died and France became a republic, I would be happy; but if France became a republic and the king of France died, I would be unhappy.

There is, in this sentence, a potential for contradiction, if only truth-conditional elements of meaning deducible from semantics are considered. However, in context the sentence is NOT contradictory, provided that the right temporal relations are assigned at the level of the utterance. In Sgaravatti and Zardini’s case, we have a potential contradiction:

(8) John knows the answer to the question how to square the circle: “In no way” but he does not know how to square the circle.

However, once pragmatic intrusion is considered, and the pragmatics overrides the semantics of ‘John knows the answer to the question how to square the circle: “In no way!” by yielding the interpretation ‘John challenged the question how to square the circle as absurd’, the contradiction which Sgaravatti and Zardini called our attention to disappears. We should not be surprised that this happens, as this is in line with many cases of pragmatic intrusion.

The case of ‘In no way!’, given as a reply to the question whether one can square the circle is clearly a case of metalinguistic negation. Metalinguistic negation, according to Horn (1985, 134) is a means for objecting to a previous utterance on any grounds whatsoever. In our case, however, the utterance is an answer to a question, not to a statement, and thus it is challenge to the (reasonableness) of the question.

1.4 Knowing how and modularity of mind.

Now, I want to address a different issue, which have a theoretical bearing on our decision to treat know how and theoretical knowledge as distinct or, otherwise, as unified. It has been claimed in the literature that knowledge how and knowledge that belong to different modules of the brain. There are studies on dissociations between knowing how and knowledge-that. (For a review, see Williams 2009; also Young 2004; for a detailed treatment of the distinction between declarative and procedural knowledge, see Devitt Forthcoming). An amnesic patient can lose her knowledge that without losing her knowledge-how. This is amazing to say the least. To those who are not persuaded, consider what happens when I am on the speedway and without seeing the car in front of me I suddenly stop my car. I am amazed. I know that I did not see that car, but I stopped my car nevertheless. It is NOT a miracle occurring or an angel that continuously assists me. We have a dual visual system, one of which is specialized in seeing objects, the other being specialized in knowledge-how. I was not involved in an accident because one visual system specialized in motion orientation was able to orient my motion path

and to block it. In other words, I know how to drive my car in a the speedway, even if I do not see the car in front of me (say, I am distracted by intense thoughts about something else).

Now, I do not want to take issue with this type of argument and claim that it is wrong. But I would like to know what it shows. Does it show that there is no connection between knowing how and knowing that? Is not an expert driver able to articulate his know-how and put that into writing, writing a manual on how to drive a car? Analogously, cannot a man who lives by buying shares on the stock exchange articulate his tacit know how and write a book on how to make investments? Assuming that there are different compartments in the mind for techniques and strategies, for skill know-how and articulated know-how, is it impossible to argue that a certain knowledge-that claim (tacit or overt) is associated with a practical mode of presentation that resides in the knowledge how module?

Surely a knowledge-how claim could be **minimal**. A person who knows how to drive the car can minimally have knowledge-that by knowing that *w* is a way of answering the question how to drive the car, while storing the practical mode of presentation with which *w* is associated in a different module of the brain. Such a minimal knowledge-that is not impossible. Stanley and Williamson's claims are not incompatible with the current theories on the modularity of mind, since they can allow that practical modes of presentation can be stored in a module which is reserved for know-how.

1.5 Knowing how and contextualism

Now I will address the problem of epistemic contextualism. Assuming that contextualism à la Keith De Rose (2009) has its validity (one can accept it fully or in part), how does the bifurcation between knowing how and knowing that fare in the light of contextualism?

Contextualism in epistemic philosophy is the claim that the same knowledge claim can have different truth conditions in different contexts. The utterance

(9) Mary knows that the bank is open on Saturday

may be true or false depending on the situation of utterance. In a low stake situation, the claim may be true, while in a high stake situation the claim may be false. Can this be applied to knowing how claims too? Intuitively one would say that contextualism provides the strongest objection to the propositional view of knowing how (How can skill cases of know-how be sensitive to low and high stakes situations?). Yet, on a pondered analysis, this is not the case. Consider the situation in which John is asked to open the safe which is connected with an explosive device and a timer. In five minutes the explosive device will explode. In this situation, if we ask whether Mary can open the safe, we may waver, even if we know that she has done that in the past and efficiently. We do not know whether she can do it in 5 minutes. The bifurcation between high stakes and low stakes situations and varying truth conditions can be transferred from knowing-that to knowing how. And this seems to support the lack of a rigid division between knowing that and knowing how.

1.6 Knowing how and Igor Douven's The Pragmatics of Belief.

Another testing bed for the assimilation of knowing how to knowing that is Igor Douven's (2010) 'The pragmatics of belief'. Igor Douven proposes that when we store a belief (in the form of an assertion or a sentence or a thought), we avoid storing it together with inferential augmentations which may lead us later to remember something which was not the case. This is called epistemic hygienics. A vivid example which comes from that paper is the reference to Gettier's problem. Suppose I know that p . Then, even if I can infer ' p or q ' from ' p ', it will not do to store in memory ' p or q ' if that is going to create trouble later, leading me to believe something that is false or unjustified. We may remember that what creates havoc in Gettier's problem is the shift from ' p ' to ' p or q '. Keeping in memory ' p or q ' when one believes ' p ' creates trouble, as that may lead to an apparently justified belief which happens to be true.

The Principle which will avoid us many problems in the future is the following:

Epistemic Hygienics (EH): Do not accept sentences that could mislead your future selves.

Other interesting examples by Douven are the following:

- (10) Peggy's car is blue;
- (11) Peggy's car is bluish.

Now, it is clear that if Peggy's car is blue, it is also bluish (blue being a stronger gradation of bluish). However, if one commits to memory 'Peggy's car is bluish' when one believes that it is blue, one will commit to memory a piece of information which will mislead one's future self. Igor Douven compares memorizing or committing to memory writing notes (e.g. Turn off the gas) which will be of use to our future selves. If memories are like notes, we should avoid writing notes that mislead our future selves.

Igor Douven's paper is of great importance to epistemology but also to pragmatics. He shows that pragmatics and epistemology are intimately connected. While Igor Douven's story can be interpreted in the light of more general principles of cognition (a memory that is misleading obviously is a case in which a believed assumption is more costly than beneficial in terms of cognitive effects; positive cognitive effects being those which put me in touch with reality, not those which drive me away from it), I cannot do this in this chapter. I confine myself to asking the innocent question whether the pragmatics of belief also applies to knowing how (knowing entailing belief). Here another crucial question arises. Does knowing how involve belief? Well, here the treatment of knowledge by Williamson may come to our aid. While all or almost all philosophers insist on the fact that belief is more basic than knowledge, Williamson (2002) insists that knowledge is more basic than belief. The case of knowing how seems to attest that he is right. One who is skilled in making a tie knot may have no beliefs on how to make a tie knot. He simply knows a practice. Of course, he must *minimally* believe that there is a way of making a tie knot.

Returning to the pragmatics of belief, it is useful to choose an example in which a strategy and not a skill is involved. Consider the case in (12), which happens to be true:

(12) John knows how to reach Trafalgar Square in 5 minutes.

Surely then he knows how to reach Trafalgar Square in 6 minutes; in 10 minutes, in an hour.

Suppose that asked whether he can answer the question whether he can reach Trafalgar Square in five minutes, John ponders on the answer that yes he can do it. Then he ponders on the logical consequences of this. He arrives at the conclusion that he can reach Trafalgar square in 10 minutes. He can also reach Trafalgar Square in 100 minutes. Suppose that, taken by the enthusiasm due to his logical deductive abilities, he goes on computing the answers to all the least stringent questions. Then he may forget the answer with which he started, which answered the question whether he could reach Trafalgar Square in five minutes. Igor Douven's epistemic hygienics will thus do some work in preventing John from memorizing all the logical consequences of his answer, which, if pursued in an endless line, will distract him from keeping his answer in mind.

So the answer is 'yes', knowing how obeys Igor Douven's epistemic hygienics. Thus the case from a dichotomy between knowing how and knowing that is made weaker.

PART II

PRAGMATIC INTRUSION

2. Pragmatic Intrusion.

Pragmatic theory (usually) deals with inferential phenomena that are intentional. On the one hand a speaker gets a message across by using non-linguistic means (or by combining linguistic means with features of the context that point interpretations in the right direction), on the other hand a hearer is busy trying to retrieve the speaker's intention. Communication must be essentially seen as a matter of projecting and recovering intentionality through the multiple clues, linguistic and otherwise, disseminated by the speaker and available to the hearer. The speaker uses a number of clues to get across her intention aware that the hearer will use them for recovering it. While Grice's project confined itself to recovering and studying conversational implicatures (inferences that are mainly cancellable, reinforceable, non-conventional and calculable), neo-Griceans (e.g. Levinson 2000) and Relevance Theorists (Sperber and Wilson 2002; Carston 1999) recognized that more is at stake than conversational implicatures and that the same processes available for the projection and recovery of conversational implicatures can be put to use in the recovery of explicatures, inferential contributions (or augmentations) that make an utterance truth-evaluable. Relevance theorists ponder on the issue of the truth-evaluability of sentential meanings, and come to the conclusion that sentential meaning is mainly underdetermined. In other words, were we to know **ONLY** what a sentence means, we would know little of what is being communicated, and the sentence would not be truth-evaluable (in many cases), simply because many other elements of a pragmatic kind enter into truth-evaluation. I have written about this in Capone (2006, 2008, 2009) by exploring the issue of whether the pragmatically obtained elements of explicatures are cancellable or not. My main conclusion is that explicatures are **NOT** cancellable. This is an assumption that can be put to use

in this chapter on knowing how and pragmatic intrusion. I will not dwell on pragmatic intrusion here, because the reader has the opportunity to read about it elsewhere; however, I will produce a number of examples which have been taken to represent cases of pragmatic intrusion. Consider the case in which Mary is asked whether she would like to join for breakfast and she replies “No, thanks; I have already had breakfast”. Clearly, she does not mean that she has had breakfast, say, on the previous day or a week ago; presumably she means that she has had breakfast on the morning of the invitation. This makes her contribution relevant to the previous invitation. Similarly, if someone asks you whether you have turned the gas off, she presumably intends to ask you NOT whether you turned the gas off ten days ago, but a few minutes ago, say before leaving the house. One of the most celebrated cases of pragmatic intrusion is the following one by Carston: ‘It is better to meet the love of your life and get married than get married and meet the love of your life’¹. Now, clearly settling temporal specifications into utterance interpretation is of importance in the evaluation of the utterance. ‘Better’ requires that the two propositions compared be different, but UNLESS we specify (by free enrichment) the temporal specifications inside each sentential fragment, we either come up with a sentence which is potentially false or with a sentence whose truth-conditions are radically underdetermined and about which we cannot say whether it is true or false. Various authors have written about pragmatic intrusion. There is no space here for a review, however see Capone 2006 for a review; also Claudia Bianchi 2001.

Pragmatic intrusion, in the case of ‘knowing how’ occurs on a heavy contextual basis. Depending on what the context is, you must choose a specific modal, a ‘de re’ interpretation or the quantifier ‘some’ (rather than ‘all’). I have argued elsewhere that explicatures (at least actual explicatures) are not cancellable. This is strong proof that the explicature is a crucial contribution to what is said, as it contributes a layer of truth-conditional meaning.

Suppose you are interested in Mary’s know-how because you want to carry out a certain job, such as opening a safe. Then it will not be indispensable to know whether Mary knows all ways to open the safe. One way is sufficient. We are not interested in exhaustive knowledge. Furthermore, in this situation, I am

¹ This example is originally Cohen’s (1971).

not interested in knowing whether Mary knows how one would open the safe, but in I am interested in knowing whether she knows some way for her to open the safe. I am not interested in knowing whether she has seen the procedure effected in the past, but in whether she has herself effected the procedure. In addition to that, I am interested in whether there is a way w such that she can open the safe with w . I am not interested in a modal such as ‘must’. In other words, I leave her free to choose one way or the other, in case she knows more ways to open the safe.

2.1 On the uses of ‘knowing how’.

What does one communicate when one says (13)?

(13) John knows how to fix the bike.

Is it different from what one says through (14)?

(14) John knows how Napoleon lost at Trafalgar.

Intuitively, one feels that (13) is connected with expertise, as John has the expertise to fix the bike, even if he may not be able to articulate his expertise apart from pointing to his actions “This is the way I do it”.

On the contrary, we have a feeling that (14) attributes to John articulated knowledge, which he may be able to articulate through a series of answers.

The question is even more complicated, as there are contexts in which the speaker’s meaning is ambivalent, pointing both to expertise and to articulated knowledge (Damschen 2009). Consider, in fact, (15)

(15) John is a driving instructor. He knows how to drive.

Unlike an ordinary driver, not only does John know the technique involved in driving cars, but he can also articulate it through answers to questions by the learners. He must be able to teach learners, not only by pointing to the way he does it (This is how you must do it!), but by giving articulate answers to questions and by explicit teaching.

So, it is clear from the start that we are faced with an ambiguity and even with cases in which both readings are possible simultaneously (the two readings not being mutually exclusive). Possibly we are faced with an **interpretative** ambiguity. Since one of the explicit aims of pragmatic theories is to avoid semantic ambiguities (see Jaszolt 1999), from the start we opt for an interpretative ambiguity, which means we must aim at providing a logical form, which, by interacting with contexts, furnishes specific interpretations.

It is very instructive to consider an asymmetry pointed out by Snowdon (2003), between knowing that and knowing how.

John knows that $p \rightarrow$ I know that p

NOT: John knows how to $p \rightarrow$ I know how to p

I take \rightarrow to be a symbol of entailment.

Actually the phenomenon should be described with a greater number of details.

A person who utters ‘John knows that p ’ is committed to knowing that John knows that p . Since that person knows the obvious consequences of what she knows, she also knows that p , given the fact that if John knows that p , p is true.

The asymmetry pointed out by Snowdon (2003) should remind us that differences exist. However, there are similarities too.

One such similarity, it should be stressed, is that ‘knowing how’ can be **communicated**. In these cases, it behaves like ‘knowing that’. Knowledge can be transmitted through speech acts, if the proposition known is not yet known to a third person and is deemed to be of use to her (say, it bears on a question which is of crucial importance to her either because she asked a question or because her behavior evinced eagerness to know a certain fact). One can say things such as:

- (16) John taught me how to drive a bike;
- (17) John learned how to ride a bike from Mary;
- (18) John told me how to cook lasagna;
- (19) John explained to me how to make the exercise;
- (20) He taught me how to cook lasagna and I remember well.

Depending on the scenarios in which these sentences are used, they may implicate that the way one learned or was taught was a verbal way. Words like ‘teach’, ‘tell’, ‘explain’ all involve the use of explicit instructions, maybe combined with showing. They strikingly contrast with ‘show’, which involves non-verbal ways of teaching. ‘Learn’ is perhaps neutral between learning by being taught through verbal instruction and learning by being shown how one does something.

There are uses of ‘knowing how’, however, which are very different or look very different from ‘knowing that’:

- (21) John knows what being operated on is like;
- (22) John remembers what falling down the stairs is like;
- (23) John imagines what falling down the stairs is like;
- (24) John tried to explain to me how it hurts to fall down the stairs.

What is known in these cases is an inner experience, which cannot be communicated or can only be communicated partially.

There are uses of ‘knowing how’ which are odd, probably because knowing how involves explicit instruction, practice, and ability.

- (25) ? John found out how to cook lasagna;
- (26) ??John guessed how to cook lasagna;
- (27) Mary knows in part how to cook lasagna.

If John was shown two cards, each describing a procedure for cooking lasagna, and just tossed a coin to guess which card is right, he would not exhibit knowledge that this is the right procedure for cooking lasagna. John’s finding out

how to cook lasagna is strange and must be interpreted as saying that John found the instructions how to cook lasagna (but he need not know how to cook lasagna, because he is not good at following instructions). (27) is used to support the distinction between knowing how and knowing that. It is claimed that one can have partial knowledge that, but not partial knowledge how. It seems to me that things are the other way round. One cannot have partial knowledge that, but one can have partial knowledge how. Consider, in fact, a worker who is instructed on how to execute a procedure; being late, he will be taught the rest of the procedure the day after. In this case, he has partial knowledge how, which will be developed into full knowledge-how when certain other skills will be integrated.

It has long been debated, since Ryle (1949), whether knowing how involves ability. Surely if I know how to ride a bike, I can ride a bike. However, the question to ask is whether this inference is an entailment or an implicature. Implicatures (in particular, potential implicatures) are defeasible, entailments are not (cases in which entailments are defeated are cases of loose usage). The philosophical literature is adamant on this. Snowdon argues convincingly that, while in many cases, if one knows how to p, one can p, this is not always the case. In particular, ability is neither a necessary nor a sufficient condition for knowing how to p (See also Bengson 2009). The standard examples discussed by Snowdon are cases like the following:

I know how to ride the bike. However, I have just had an accident and I have broken my leg. Thus, it appears that I cannot ride the bike. Nevertheless, I still know how to ride the bike. When I recover, I will ride the bike again without any problems.

Various authors recognize that while ability is neither necessary nor sufficient for knowing how, in many cases when one says:

(28) John knows how to cook spaghetti

it is implied that he can cook spaghetti.

It has been proposed that this is a conversational implicature. Surely this must be the case, given that conversational implicatures are non-monotonic inferences, and may arise standardly (standard implicatures). It is our world knowledge that licenses the ‘ability’ implicature when one says ‘X knows how to p’. We know that, if there are no special impediments, knowing how *ipso facto* creates the ability to do p. In other words, we are empowered to do p by our knowledge on how to p. The implicature that one can do things becomes an explicature when knowing how has to be put to use. Suppose I am interested in whether John knows how to open the safe. Then I am not only interested in knowing whether John knows a procedure but whether he has ever applied that procedure in the past. Depending on the context, the ability explicature becomes consolidated, and it may even become hard to cancel.

2.2 Rumfitt and ‘savoir + infinitive’ : towards a pragmatic story

Rumfitt (2003) proposes a strong objection to Stanley and Williamson. He considers a wide range of languages in which ‘knowing how’ is expressed through a construction of the type ‘savoir + infinitive’ (e.g. French, Italian, Russian). It appears that there is a difference between

(29) Giovanni sa nuotare
(lit. John knows to swim)

and

(30) ?Giovanni sa come nuotare.
(John knows how to swim).

If well-formed at all, (30) seems to conform to Stanley and Williamson’s semantic elucidation for knowing how (There is a special way *w* of swimming and Giovanni knows *w*). (30) is probably appropriate to situations in which one, out of

a several, ways of swimming is suited to the situation (in that a particular goal is being pursued). One also has the feeling that (30) allows a continuation as in:

(31) John knows how to swim (to impress the examiners).

The materials in parentheses are supplied contextually through free enrichment. So summing up, one has the impression that ‘sapere nuotare’ and ‘sapere come nuotare’ have different meanings.

Stanley and Williamson can reply to Rumfitt that ‘savoir + infinitive’, after all, is a ‘know how’ construction in so far as it is possible to analyse it as having a question word as part of pragmatic intrusion (free enrichment, presumably). But then the question arises why there should be such a marked difference between (29) and (30). A possible answer to this question – an answer which is favorable to Stanley and Williamson – is that there is surely a contrast between (29) and (30), but that is NOT a matter of semantic entailments, but a matter of explicature. Assuming that (29) and (30) have the same semantic/syntactic structure, the more prolix and more marked (30) triggers some explicatures (M-explicatures, according to Horn (2009), Levinson (2000), and Huang (2000)). According to a Relevance Theory perspective, given that the more explicit ‘come’ would be associated with the same semantic readings, it would be more costly to the hearer to utter one more word; thus, to make this cost bearable, one would have to offset the processing cost with suitable contextual effects. Thus ‘come’ triggers a conversational inference which aims at completing the sentential structure with a purpose clause.

2.3 Inferences in context

When one examines the cases of knowing how, one may notice some differences among the following cases.

- (32) John knows how to drive the car;
- (33) John knows how to fix the bike;
- (34) John knows how to drive the bike;
- (35) John knows how to cook lasagna;

- (36) John knows how to persuade the coach to let him play;
- (37) John knows how to teach a history lesson;
- (38) John knows how to arrive to the office by the shortest route.

All these cases are similar; yet they involve different shades of meaning. In (32) we seem to be confronted with a case of (skill) know-how. However, if the sentence is embedded in a context in which a mother tells her son to ask for John's help in learning to drive the car, it is conversationally implied that her son should be ready to provide verbal instructions. (33) seems to implicate that John also knows how to find the tools needed to fix the bike. (34) seems to express a skill; however, if the sentence is situated in a context in which you are asking for help on how to use a professional bike, 'bike' here is interpreted as 'professional bike' and the sentence is promoted from describing a skill to describing competence which can be articulated through instructions. (35) may range between an interpretation according to which John is being attributed a skill (he knows how to cook lasagna by route) and an interpretation according to which he is being attributed a source of knowledge (He knows how to cook lasagna; he has a cookery book in the kitchen and he resorts to that book when he needs to cook lasagna. Here 'knows how' actually implicates, perhaps due to a loose usage, that one knows 'where' to find a recipe). (36) implicates that John has worked out a strategy on how to persuade his coach to let him play. (37) may implicate that John has a special method for teaching a history lesson or that he has rehearsed a history lesson. (38) implicates that John is able to calculate the shortest route to his office. It is incredible that, in different contexts, 'knowing how' sentences take on different shades of meaning, which are not expressed by semantics alone.

There are other types of pragmatic interpretations with knowing how sentences. Consider, for example, scope ambiguities, as in the following:

- (39) John knows how to fix a bike;
- (40) John knows how to ride a bike;
- (41) John knows how to calm up the dog;
- (42) John knows how to fix every bike.

In (39) and (40) it is unlikely that the indefinite NP has scope over the verb of propositional attitude. Thus we exclude the interpretations:

- (43) There is a bike: John knows how to fix it.
 (44) There is a bike: John knows how to ride it.

The natural interpretations of (39) and (40) seem to be the following:

- (45) If there is a broken bike, John knows how to fix it.
 (46) If there is a bike, John knows how to ride it.

In (41), ‘the dog’ escapes the scope of the verb of propositional attitude and has scope over it. As Jaszczolt (1999) correctly insists, NPs tend to be correlated with ‘de re’ readings by default. I have argued in Capone (2009) that this story must be coupled with a RT explanation. I cannot go into this here. (42) has different interpretations. If there is a domain of quantification which is salient, then ‘every’ will be interpreted as quantifying over elements from that domain (See Stanley 2007). If there is no such a domain present, then it will have a genuinely universal quantification reading. HOWEVER, it will not be understood that John knows how to fix every bike (at the same time). What is probably needed is a selection function that selects the variables individually, and NOT cumulatively.

For all x , x a bike, John knows how to fix it (SEL FUNC one by one);

For all x , x a bike, John knows how to fix them (SEL FUNC all at once).

Presumably bound variables must have formal properties which correlate with different selection functions.

World knowledge is what enables the speaker/hearer to have access to the most reasonable or plausible reading. It is implausible that one can fix all bikes at the same time. Thus this interpretation is NOT selected. The default interpretation is the one that accords with socio-cultural defaults (see Jaszczolt 2005).

There are other interpretation problems concerning ‘knowing how’ interpretations. Let us assume that part of the story of knowing how utterances is that they are interpretable in the following way:

(47) John knows how to ride the bike

(48) John knows that there is a way of driving the bike *w* and this is a way for him/for one to drive the bike.

The choice between one interpretation or the other is not trivial, and obviously depends on the context of utterance. A certain context of utterance will favor an interpretation rather than the other. However, one gets the impression that the default reading is (49):

(49) John knows that there is a way of driving the bike *w* and this is a way for him to drive the bike.

Presumably, the reading which has greater contextual implications (cognitive effects) wins. While in the majority of cases, there is a coincidence between the way one drives the bike and the way John drives the bike, we are describing the situation from John’s perspective, not from everyone else’s perspective; thus it is natural that John has developed a way to drive the bike that accords with his natural inclinations. This is clearer in the following example:

(50) Alessandro knows how to drive the car in an Italian speedway.

Italian speedways are such that when there is a car from an incoming road, you either need to slow down in order to avoid accidents or you drive in the left lane, in order to avoid an accident. Surely this is an idiosyncratic way of driving. But Alessandro has developed his personalized way of driving. I assume that the natural interpretation of (50) is that Alessandro has his own way of driving in an Italian speed way. This is the interpretation which has richer cognitive effects and the one which conforms to Alessandro’s perspective.

One further problem is the interpretation of the deontic modal.

(51) John knows how to fix the bike

is equivalent to

(52) John knows that there is a way w and that is the way to fix the bike.

(52) Roughly means (53):

(53) John knows that there is a way w and w is the way he can fix the bike,

rather than

(54) John knows that there is a way w and w is the way he must fix the bike.

It is a matter of pragmatics that ‘can’, rather than ‘must’ is chosen. ‘Can’ involves a greater space of freedom than ‘must’. One assumes that activities like fixing bikes, etc. are done according to free will, and not under coercion. Again it is socio-pragmatic defaults that are at work here.

Another interesting interpretative problem is whether one should have one of the following interpretations of say (55).

(55) John knows how to arrive at Trafalgar Square in 5 minutes.

(56) John knows some way w and w is a way he can reach Trafalgar Square in 5 minutes;

(57) John knows all ways w such that John can reach Trafalgar Square through w .

It appears philosophers prefer the interpretation (56). Again socio-cultural defaults are at work here. It is rarely the case that one knows all ways to do p and to know how to do p , it suffices to know one way to do p , rather than all. The most stringent interpretation is thus discarded. One settles on ‘some’ as a default.

There are other inferences to consider. There are implicit restrictions to knowing how claims. Consider, in fact, the following (from Hawley 2003):

- (58) John knows how to drive (American cars);
- (59) She knows how to drive (European cars);

- (60) She knows how to cook (for an informal dinner);

- (61) She knows how to cook (as a chef in a restaurant).

In all these cases, the context of utterance restricts the truth-conditions of the utterance, Contextualism can provide an explanation for why there can be so much variation in the truth-conditional import of the utterance. In all of the cases above, there is a restriction on the truth-conditional claim; this restriction is implicit and must be inferred pragmatically. Typical scenarios and situations of use can provide the materials for the pragmatic restrictions. Of course, one might reply (à la Cappelen and Lepore 2005) that one must then allow for an unlimited number of pragmatic increments. (58) could mean:

- (62) John knows how to drive in a Sicilian speed way;
- (63) John knows how to drive in a chaotic traffic;
- (64) John knows how to drive among drunk people;
- (65) John knows how to drive a scooter.

I agree with Cappelen and Lepore that there can be many (even too many) contextual increments. But this does not make contextualism less appealing, provided that we know how to combine the basic sentential meanings with the contextual variables that lead to possible increments.

3. A unified treatment of knowing how: pragmatic intrusion?

It has been claimed that knowing how utterances are characterized by an interpretative ambiguity. They are sometimes understood as describing abilities or

skills; they are sometimes taken as describing knowledge of ways one can do p. Surely, there are a number of cases in which ‘knowing how’ sentences appear to describe skills. The blacksmith has mastered a certain skill. He need not be able to articulate his knowledge, since many of the operations he executes are automatic. However, John who knows how to reach Trafalgar Square in 5 minutes does not have a skill. He was able to perform a calculation and to work out the shortest route to Trafalgar Square. He is able to say that w is the way he knows to arrive at Trafalgar Square in five minutes. He is also able to say why he chose w rather than w’. This is not a skill or an ability. This is real knowledge.

Given this clear bifurcation in the uses of knowing how, one can hypothesize that the different uses are based on a single logical form²:

(66) John knows ... how to p

which is then pragmatically enriched to

(67) John knows the answer to the question how to p;

(68) John knows the technique how to p.

This amounts to claiming that (66) is a fragmentary semantic schema, to be enriched in context. Certain contexts promote skill interpretations; other contexts promote knowing the answer to the question whether p interpretations.

Another solution might be to follow Stanley and Williamson up to a certain point. At some point Stanley and Williamson claim that the way in which one can p is presented to the speaker/hearer through a practical mode of presentation. But this is clearly suitable to skills, not to know-how that is theoretical. Furthermore, we should distinguish between a practical and an indexical mode of presentation. The indexical mode of presentation is of use in cases like the following:

² Brogaard (2009) leads us in this direction, although he is not specific about pragmatic intrusion and what kind of underdetermined logical form is needed.

(69) John showed me how to fix the bike and I learned it.

For this example to work, John, rather than explaining how to do things, showed the hearer how to fix the bike, presumably saying “This is the way to do it”. This is not only a practical mode of presentation, but an indexical mode of presentation.

Now, I am favorable to having a unitary logical form like:

(70) John knows the answer to the question how to fix the bike.

One can in fact enrich this logical form by making use of a practical mode of presentation. This surely is more compatible with the situation in which both a practical and a theoretical mode of presentation is involved. Consider again the situation in which one says

(71) John knows how to drive the car,

implying that, since John is a driving instructor, he has both articulate knowledge which he can teach and practical knowledge which he can show. In this case, it is not clear given the logical form

(72) John knows how to drive the car

how to arrive at the multiple interpretation (skill and articulate knowledge). Surely (73) is a bit far-fetched:

(73) John knows the answer to the question how to drive and the technique how to drive.

However, (74) is not:

(74) John knows that there is a way w in which he can drive the car.

W can be presented through a practical or theoretical mode of presentation; but there is nothing to prevent the speaker from using and expecting the hearer to infer a double mode of presentation of this way of driving.

Why is (73) far-fetched and implausible? I assume that a reason why (73) is implausible is that it says that John knows two (distinct) things, and this does not accord with our understanding of daily use, while (73) can be interpreted as meaning that John knows one thing, which may be presented to him in different ways, according to the demands of context.

3.1 A comparison between the pragmatics of knowing how and the pragmatics of propositional attitude reports.

Ideally, we would expect the pragmatics of knowing how and the pragmatics of propositional attitude reports to behave similarly. Let us remind our readers of my treatment of the pragmatics of belief reports in Capone (2008). There I established some hidden constituents as a result of free enrichments. In particular, I posited some null appositives, attached to the that clause of belief reports. Schematically, this amounts to the following:

(75) John believes that [Mary is in Paris] [Mary is in Paris] APP

The apposition, attached to the sentential constituent of the that-clause, makes it necessary that the semantic referents associated with ‘Mary’ and with ‘Paris’ in the structure of the thought or belief attributed to John should be glued with phonetic forms such as ‘Mary’ and ‘Paris’, which explains why substitution of identicals is blocked inside belief reports. In Capone (2008) I exposed a pragmatic story according to which the principle of relevance is involved in enriching the structure of (75) by adding an apposition clause. Such a story was based on cognitive effort, and assumed that, if a speaker bothered to use names in order to talk about referents in reporting a thought of John, then he had a reason to do so, this reason being that the names figured as parts of John’s thoughts. Whether this

story has to be amended or amplified is another matter, which I will not investigate here.

The question I am interested in is whether the pragmatics of knowing how works in a similar way. And the answer is that it doesn't exactly work in the same way, attributing to John knowledge of a way of doing an action. In fact, we have not said anything on whether the way of doing the action is outside or inside the scope of 'know'. One possibility is that it is outside the scope of 'know'. We, thus, accept the idea that, on a given interpretation, an utterance of 'John knows how to ride the bike' amounts to:

(76) There is a way w , such that John knows how to ride the bike in way w .

Now, this logical form explains why, on some occasions, knowing how utterances are interpreted in a transparent way. Of course, the possibility remains open that the quantifier in 'There is a w ...' can occur inside the scope of 'John knows that' and, in this case, opacity is created, which justifies the considerations of those who posited opacity as part of knowing how structures. But now, if there is an ambiguity between a quantifier being outside the scope of 'know' and a quantifier being inside the scope of 'know', it is clear that my pragmatic story works well when the opaque reading is chosen, but NOT so well when the transparent reading is chosen. I suggest that we go for a bifurcation. When the opaque reading is chosen, then in the same way in which belief reports could be equipped with null appositions through pragmatic intrusion (specifically free enrichment), a null apposition can be attached to w , and this apposition would amount to a mode of presentation of w . We assume that a practical mode of presentation of w is provided by default, given that a practical mode of presentation is often a firmer way of teaching one how to execute a procedure, since showing a procedure provides all or most of the information which could be provided through a linguistic articulation of the way in which something can be done, but things do not work all the way round. In other words, it is not the case that a linguistic articulation of the way something can be done will be able to express all the

information which one can furnish by a practical mode of presentation (by giving an example of the action).

But if this story works for the opaque reading, this story should also work for the *de re* reading, since even outside the scope of a verb of knowledge, a way of doing something can be enriched by free enrichment through an apposition structure and a practical mode of presentation can be chosen on the grounds that it is necessarily more informative than a linguistically articulated mode of presentation.

We have seen that the generalized implicature story is only part of the full pragmatic story. Generalized implicatures must, in fact, be tailored to specific contexts and, therefore, it is not unusual that a certain implicature may be aborted in favor of an alternative interpretation. Examples (32) – (38) seem to prove that contextual considerations must be accommodated when one enriches the logical form of a specific knowing how utterance.

So we have abandoned the simplistic assumption that generalized implicatures are selected despite the context or regardless of the context of utterance. We only retain the weaker view that generalized implicatures are triggered by default and are promoted to actual interpretation only if the context allows them (in other words if the context filters them in). If the context filters them out, then interpretation must follow a procedure in which various steps of reasoning are involved in selecting the favored interpretation.

Conclusion

Devitt (Forthcoming) argues that epistemology should be divorced from linguistic considerations as Stanley and Williamson's considerations are mainly based on linguistic considerations, but NOT on consideration of metaphysical issues. Now, while surely there are reasons for Devitt's assertion, this chapter in the main shows that philosophy and pragmalinguistic considerations are deeply entangled. When matters of inference are considered, it is also worthwhile considering pragmatic intrusion. There are reasons to believe that almost all questions that

pertain to knowing how can be analyzed with reference to pragmatic intrusion – that being a pervasive feature of human communication. Given what Igor Douven says about the pragmatics of belief, we are not surprised that knowledge, as well as belief, is strictly connected with the notion of communication and that many of the arguments that pertain to the sphere of epistemology can be illuminated by the tool of pragmatic intrusion.

Chapter 10

The semantics and pragmatics of attitudes 'de se'.

Abstract.

After reviewing a number of theories of attitudes 'de se', in particular philosophical and linguistic theories, I adopt the theory by Higginbotham and, after showing that there are problems in it, I complement it with an theory of pragmatic intrusion, which supplied an Ego concept to the logical form of the 'de se' attitude. I also articulate the internal dimension of attitudes 'de se' by pragmatic intrusion, whether partial or full.

It is widely accepted that there is a huge gap between the meaning of a sentence and the messages actually conveyed by the uttering of that sentence (Huang 2007).

But why do we need clear, sharply demarcated boundaries at all, when pragmatics is in constant development, so that boundary markers, once placed, will have to be moved all the time? Maybe a 'pragmatic definition of pragmatics could be found that avoids both the Scylla and Charybdis of the above alternative? (Mey 2001, 7).

Introduction

Since Mill (1867), there is a convergence in pragmatics on the idea that "What a speaker intends to communicate is characteristically far richer than what she directly expresses; linguistic meaning radically underdetermines the message conveyed and understood. Speaker S tacitly exploits pragmatic principles to bridge this gap and counts on hearer H to invoke the same principles for the purposes of utterance interpretation" (Horn 2004). In the spirit of Horn's programmatic assertion, in this chapter I shall deal with utterances reporting attitudes 'de se' and I shall suggest that pragmatic principles and mechanisms are at play in the recovery of 'de se' attitudes in English (and in Italian), in those cases in which grammatical and semantic information does not exclusively determine the 'de se' interpretation. The examples I shall consider are cases like 'John believes he is clever' or 'John remembers walking in Oxford', where a subject has a thought or a memory about himself/herself (the subject of the thought) and these are contrasted with examples like 'John believes he is clever' ('He' used here to refer to someone other than John, whom the speaker is pointing to), in which the pronominal has a 'de re' interpretation. In cases like 'John remembers walking in Oxford' I assume, following Higginbotham (2003) that PRO (the null subject of the subordinate clause) is associated with a 'de se' interpretation).

I shall argue that 'de se' attitude attributions are not completely reducible to pragmatic mechanisms that take input from 'de re' logical forms. I shall point out what problems radical pragmatic analysis has to face. I propose that the pragmatic processes at work in those cases where semantics and grammar do not fully determine the 'de se' reading are cases of free enrichment (John believes he* is

clever). In free enrichments, although there does not seem to be either an overt indexical or a covert slot in the linguistically decoded logical form of the sentence uttered, the logical form nevertheless needs to be conceptually enriched in the explicature (Huang 2007, 191). I cast my view in the relevance-theoretic framework, according to which the human cognitive system works in such a way as to tend to maximise relevance with respect to communication, the principle of relevance being responsible for both the explicit and implicit content of an utterance. However, it should also be said that I was greatly influenced by Levinson's (2000) ideas about pragmatic intrusion. I propose to add another type of intrusive constructions: sentences expressing attitudes 'de se'.

The chapter has the following structure. I initially present philosophical theories of 'de se' attitudes, starting from Castañeda (1966), distinguishing propositional treatments from property-theoretical positions. I then discuss a linguistic treatment of 'de se' attitudes' (Higginbotham 2003) as well as some important objections by Safir (2004). In the final sections I discuss pragmatic intrusion and I distinguish cases in which 'de se' inferences are cancellable and those where they are not. I propose that PRO is associated with 'de se' beliefs through semantics alone (pragmatics is not involved), but that the internal dimension of PRO is a pragmatic inference. I also propose that Higginbotham's treatment does not do justice to the fact that 'de se' beliefs involve a mode of presentation that incorporates a mental tokening of 'I'. This component of meaning is on top of Higginbotham's semantics for PRO first-personal readings. (I modify his semantic elucidations to accommodate this intuition). As Recanati (2007) says "Indeed, the ability to entertain implicit *de se* thoughts is arguably a necessary condition for anyone to evolve the concept EGO (Recanati 2007, 177). Analogously, I propose that there can be no 'de se' thought without a mode of presentation that somehow vocalizes the word 'I' (albeit in thought) or some transformation of it. My proposal is that the content of this word 'I' is mixture of a demonstrative and of a proper name and has to supplement Higginbotham's treatment of the first-personal reading of PRO.

To be even more schematic, I claim that 'de se' interpretations associated with the pronominal 'he' (those cases for which Castañeda uses the asterisk) are pragmatic and I explain how pragmatics predicts them. Through pragmatics, I also predict a number of 'de re' interpretations. The external first-personal dimension of PRO in

constructions such as ‘John remembers going to the cinema’ is semantically ‘de se’, following Higginbotham (2003).

Hornian, Levinsonian and relevance theories bear on the present analysis in that (1) they assume that there is a gap between the sentential meaning and utterance meaning, as is also attested in the case of ‘de se’ intrusive constructions; 2) they provide the principles of the analysis of the pragmatic inferences which in some, but not in all cases, are responsible for ‘de se’ interpretations. 3) These theories are compatible with Jaszczolt’s (2005) theory of merger representations, representations in which semantic, pragmatic and socio-cultural information merge to produce propositional forms that are truth-evaluable.

In my chapter I consider the following list of factors for the purpose of my pragmatic analysis:

List of factors in ‘de se’ attitude analysis

Expressions	Verbs	‘de se’ or ‘de re’	Internal	External
PRO He, him His himself He himself	Believe remember imagine Expect Dream Forget Know...	de se de re	Semantic or Pragmatic	Semantic or Pragmatic

In particular, I argue that ‘imagine’, and ‘expect’, ‘dream’, ‘know’ are different from ‘remember’ in terms of the internal dimension.

1. Philosophical perspectives

Although my treatment is a linguistic one, in that it mainly deals with interpretations of utterances and with a systematic exploration of minimal pairs, there is no denying that the topic originated in philosophy. I therefore start this chapter with an orderly

presentation of theories of ‘de se’ beliefs and other (propositional) attitudes. The verbs we are going to analyse are verbs such as ‘believe’, ‘remember’, ‘image’, ‘want’, etc., which are usually referred to as verbs of propositional attitude. Since ‘de se’ beliefs specifically raise doubts as to whether propositions are involved in the analysis of such verbs, it is best to use the neutral term ‘attitudes’. The theory of ‘de se’ attitudes is clearly a topic within the philosophy of mind; however, here I shall be more narrowly concerned with the linguistic implications of the philosophical theories. In other words, I shall be mainly preoccupied with matters of inference and, specifically, of linguistic entailment. It is not possible, in this section, to review in great detail all the articles or books on ‘de se’ beliefs – and I am sure there are ramifications we have to explore in the future. However, it is possible to follow the ideas on ‘de se’ interpretations as triggered by Castañeda’s seminal article which is the basis for my pragmatic interpretation of ‘de se’ attitudes, through the development by Perry who insists on the causal efficacy of having ‘de se’ thoughts. The discussion of Millikan and Shoemaker is essential to my claim that ‘de se’ readings involve the use in thought of an indexical like ‘I’. Harcourt explores the view that ‘I’ could be an implicated component of ‘de se’ readings. This too connects with my idea that ‘de se’ readings involve the concept of ‘I’. I also discuss the contributions by Boër & Lycan (1980), who oppose the essential indexical story by Perry, and I defend the story by Perry. Finally, I discuss the property-based view of belief, as an alternative to propositional theories of the attitudes.

1.1 Castañeda on ‘de se’ beliefs.

According to Castañeda (1966), there is a difference between (1) and (2)

- (1) The editor of Soul believes that the editor of Soul is a millionaire;
- (2) The editor of Soul believes that he* is a millionaire.

Specifically (1) can be true, without its being the case that (2) is true. Suppose that John has been informed of the fact that the Editor of Soul has inherited a huge sum of money: then he knows that the editor of Soul is a millionaire. However, he has not been informed of a sudden change in the board of Soul and, specifically, he does not know that he himself has been appointed editor of Soul. Then he does not know that he himself is the editor of Soul, albeit he knows that the editor of Soul is a

millionaire. Since the pronominal 'he' is ambiguous between a 'de se' and a 'de re' interpretation, Castañeda uses the asterisk to disambiguate. The asterisk will turn the pronominal into an **essential indexical** (presumably it is these asterisks which are the topic of our pragmatic analysis, a linguistic fact neglected or not brought into focus by the philosopher and his followers).

Presumably essential indexicals complicate theories. Consider what happens to the following considerations from the theory of knowledge. If I know (3)

(3) John knows that Paris is not in Great Britain,

I can infer:

(4) Paris is not in Great Britain

and I can assert (4) (The proposition expressed by (4) can be detached from (3)). However, things become more complicated when we take sentences such as (5) into consideration:

(5) John knows that he* is clever.

We simply cannot infer

(6) He* is clever,

Albeit we can infer (7)

(7) He is clever.

If I were to infer (6), I would have to think of myself that I am clever. Of course, it should be said (and Castañeda does not say this) that the interpretation (6) is obtained by recourse to pragmatic principles and that the inferential step from (5) to (6) needs to be mediated by pragmatics – thus it is out of the question that (6) is analytic in (5).

1.2 John Perry (1979) on ‘de se’ beliefs.

Perry’s story develops the considerations by Castañeda, by linking the ‘de se’ notion to the theory of action, claiming that the ‘de se’ concept is causally active. Perry (1979) holds a line of thought similar to Castañeda’s. His story about the supermarket is an impressive attempt to connect the issue of belief (and, in particular, de se beliefs) with the theory of action. John Perry is in a supermarket and sees a trail of sugar left by what he thinks is a different shopper. He follows the trail of sugar because he wants to tell the unaware shopper about it, until it dawns upon him that he (himself) is the messy shopper. He stops following the messy shopper when he understands that he himself is that person. The belief that the messy shopper is leaving a trail of sugar in the supermarket is not causally relevant to taking action, instead the belief that he himself is leaving a trail of sugar will prompt him to take action. The mode of presentation involved in the belief state is, thus, causally involved in determining a certain action and it is important that a first-personal mode of presentation causes an action which would not be caused by a non-first-personal mode of presentation. The account presented so far is, in principle, compatible with Castañeda’s considerations.

Perry (2000) differs significantly from Castañeda’s ideas, though. Perry focuses on the pragmatic nature of the inference involved in a sentence such as (8)

(8) Privatus believes that he(*) is rich.

According to him, that ‘he*’ is an essential indexical is due to a pragmatic process (Perry does not bother to explain the details of this process) – in fact, the inference can be cancelled. Suppose that Privatus is acting in a play and that a speaker utters (8) meaning that Privatus believes that the character he is acting out is rich. It follows that Privatus does not believe himself to be rich. Hence the interpretation of ‘he*’ as an essential indexical is not a semantic one (and it is optional, as indicated by (*)). Since the interpretation is due to a pragmatic process, it can be cancelled (in this case it is cancelled by contextual assumptions).

Furthermore, Perry believes that even (9) is only pragmatically ‘de se’:

(9) The dean was surprised to find that he believed himself to be overpaid.

Such a sentence can be uttered in a context in which the dean has complained that professors who publish less than ten articles per year are overpaid. (It just happens that the dean has published less than ten papers per year, but does not remember that). It is clear in this context that (9) does not mean that the dean believes himself to be overpaid. Perry reflects on some interesting cases of cancellability, thus paving the way for a pragmatic theory of belief reports (the question whether he is simply making use of parasitic or etiolated cases of language use is not important for the time being; I hope to be able to place the pragmatic theory on a more solid footing).

1.3 Millikan

Unlike Perry and others, Millikan (1990) proposes that essential indexicals are different from ordinary deictic expressions. Millikan, unlike Perry, believes that deictic expressions have nothing to do with action. In fact, only in the case in which the deictic expression identifies the first person perspective in action, is action influenced by the deictic expression, but this is the case in which the deictic expression is nothing but a mode of presentation of the ego. The author believes that there is a noteworthy difference between ordinary deictic expressions and the essential indexical. Ordinary deictic expressions have their referents identified through the context of utterance. Instead, an essential indexical is necessarily related to the first person perspective, as the thinking subject directly presents himself to conscience. For Millikan it is reasonable to use a mode of presentation (e.g. @RM) similar to definite descriptions or proper names except for the fact that its use implies the identity $I = @RM$. This mode of presentation is connected with dispositions to act and, in this sense, is causally active.

1.4 Shoemaker (1968).

Shoemaker's considerations intersect with Millikan's. On the one hand, in fact, Shoemaker ponders on the similarity between 'de se' concepts and indexicals, on the other hand he speculates that even indexicals require some identification through perception, and perception can go wrong. Thus, we can find a link between Shoemaker's notion of immunity to error through misidentification and Millikan's notion of the essential indexical as a special concept, similar to a proper name. Shoemaker, unlike Millikan, finds a similarity between the essential indexical and the use of deictic expressions. In fact, he claims that in certain cases (albeit not in all

cases) the use of a demonstrative serves to identify an object and it is simply impossible not to identify the object demonstrated by the speaker in context. If I say ‘Look at this dog’, there is no chance that the object is not identified or not identified correctly. Of course, Shoemaker is aware of cases in which the use of the demonstrative gives rise to misidentification, as when I say ‘This is a red necktie’ and I identify the necktie because I touch it and identify the colour through tactile sensations, being aware that all my silk ties are red. Immunity to error through misidentification arises with psychological predicates, as when I say ‘I am in pain’, albeit not with non-psychological predicates that involve reference to parts of the body, as these parts could be easily misidentified. Like Millikan, Shoemaker believes that the question of identifying the ‘I’ does not arise and that the ego is available to consciousness directly, through a mode of presentation. (For a deeper treatment of the issue of error through misidentification, dealing with the possibility that a memory that belongs to another person is implanted into my mind, see Evans (1969), Prior (1998) and Coliva (2006)).

1.5 Harcourt (1999).

Harcourt too believes that essential indexicals have a first-personal interpretation and resorts to a conventional implicature analysis. Harcourt makes use of a Davidsonian theory of propositional attitudes and believes that it is useful to analyse e.g. Mario believes that Joan is in Paris as (10)

(10) Mario believes that: Joan is in Paris.

The crucial problem for Harcourt is to explain how first-personal modes of presentation interact with the theory of action, while preserving semantic innocence. In fact, changing the example, and using a ‘de se’ case, the problem is how to relate

(10) Mario believes that he* is happy

(11) Mario believes that: he is happy

in such a way as to preserve semantic innocence. It is interesting that Harcourt discards one move that is available to him, that of conversational implicature, If he resorted to this move, he could explain how, despite the fact that (10) can be

analysed in terms of purely extensional semantics, a first-personal perspective is conversationally implicated through the usage of a logical form such as (11). Harcourt gives up the implicature hypothesis, because, in his opinion, it is not possible to test the hypothesis due to the ambiguity of the sentence and because according to interpretationism (which is the view he accepts), the interpretation of the embedded sentence requires that the first-personal interpretation be a semantic component of the content of the embedded sentence. However, it should be said that all interpretationism requires is that the embedded sentence be semantically interpreted as in the original utterance, (11), pragmatic increments being on top of that. The question of the ambiguity is easily resolved by resorting to Modified Occam's razor, which enjoins us not to multiple senses if simpler hypotheses can be considered (see Grice 1989, Jaszczolt 1999, Ariel 2008).

Harcourt believes that the essential indexical implies an original context of use in which the thinking subject presented himself as 'I' (I take he is invoking the notion of conventional implicature) – however, it is difficult to see how this treatment can preserve semantic innocence, given that only the character of the expression 'he*' in (10) can guarantee such an implication. Harcourt's theory, instead, works much better in case he is willing to defend a conversational implicature analysis.

So far we have seen cases where philosophers invoked a special 'de se' concept, which was said to be causally active. However, not all philosophers agree with this, and in the next two sections we shall discuss two groups of authors who defend the opposite thesis.

1.6 Boër & Lycan.

According to Boër & Lycan (1980), the notion of belief 'de se' is reducible to the notion of 'belief de re'. According to them a sentence such as (12)

(12) John believes that he himself is in danger

has the following semantic content:

John₁ believes that John₁ is in danger.

According to the authors, the usage of a proper name can receive two interpretations: 1) attributive; 2) referential. If the interpretation is attributive, then the use of a proper name within a belief attribution will render the belief 'de dicto'. Consider

(13) John believes that John is in danger.

John can accept 'John believes he is in danger' without having to accept (13), in that John may have forgotten his name. However, the non-equivalence of (13) and 'John believes that he is in danger' is demonstrated only if one confines oneself to the attributive interpretation of the proper name. Instead, if we opt for a 'de re' interpretation of the proper name, there is an equivalence between (13) and (12).

The article is interesting because it moves towards a pragmatic theory of 'de se' beliefs, albeit one based on conventional implicature and not one based on conversational implicature. According to Böer & Lycan, a speaker who utters (14)

(14) John believes that he himself is in danger

conventionally implicates that John must be willing to use the sentence 'I am in danger' in 'foro interno' (inside his mind). Although we are not exactly facing a case of conversational implicature, we must say, nevertheless, that this article is a step in the right direction, as we would like to maintain that pragmatics is largely responsible for a number of cases of 'de se' inferences. (See Davidson 1985 for further discussion of Böer & Lycan's work).

1.7 Davidson (1985).

Davidson's approach, like the one by Böer & Lycan (1980), also attempts to reduce 'de se' belief to 'de re' belief. He does that by saying that for a person X, an individual a and a property P, 'de se' beliefs are exactly those beliefs 'de re' where a = X. In other words, a belief 'de se' is nothing but a belief 'de re' about oneself. An obstacle to this approach is the kind of considerations invoked by Perry to distinguish between 'de se' and 'de re' beliefs. Perry claims that 'de se' beliefs cannot be reduced to beliefs in which a definite expression appears instead of the essential indexical. To avoid this complication, Davidson discusses an example by Böer &

Lycan. Consider the situation in which Van and Wilfrid are different ways of referring to the same person. John, however, believes that Van and Wilfrid are different persons (due to the presence of a mirror). He sees a wild animal about to pounce on Van. Now consider the following sentences:

(15) John believes of Van that he is in danger;

(16) John believes of Wilfrid that he is in danger.

The replacement of one name with another will change the dispositions to act on the part of John. John is willing to act to protect Van from the danger, but not to protect Wilfrid (as John does not believe that Wilfrid is in danger). So Davidson and Böer & Lycan are willing to conclude that there is nothing special about special indexicals, since analogous considerations are applicable to 'de re' beliefs. Hence there is no objection to reducing 'de se' beliefs to 'de re' beliefs.

What should we make of this type of objections raised by Böer & Lycan (1980) and Davidson (1985)? All that the authors have shown is that modes of presentation of NPs can be causally efficacious. This is hardly surprising. If I were to discover that my friend Liz is nothing other than Queen Elisabeth I would treat her more kindly and formally than I would be inclined to do under my present state of knowledge. If modes of presentation can be causally efficacious, then it is arguable à la John Perry that 'de se' modes of presentations are causally efficacious in ways that 'de re' readings are not. For Davidson's story to work, he would have to show that 'de se' readings can be reduced to 'de re' readings. In other words, he would have to show that 'de se' readings come for free through pragmatics once we have 'de re' interpretations. However, I argue in a later section that even pragmatic 'de se' interpretations require first-personal readings as marked through a syntactic construction that contains PRO at the level of the logical form of the merger representation. In other words, it is not possible to reduce in all cases 'de se' readings to 'de re' readings. Of course, a 'de se' reading may be compatible with a 'de re' reading, this is why Davidson's story seems to work. I may look at a mirror and while I think that I am clean and tidy, I also think that he (the person in the mirror) is clean and tidy. However, in case I have to shave to become clean and tidy, I will

shave myself and not the person in the mirror, because the 'de se' thought is causally efficacious, while the accompanying (compatible) 'de re' thought is not.

1.8 Property-theoretical views of belief

So far we have reviewed philosophical theories on 'de se' attitudes which were mainly propositional. However, it should not be taken for granted that all philosophers consider attitudes 'attitudes to propositions'. Some eminent philosophers saw some merit in expressing beliefs and other attitudes through properties. An important, albeit not uncontroversial view about belief is the one that considers belief basically a relation of ascription of properties. Thus, if I believe that I am happy, this amounts to ascribing myself the property 'being happy'. A belief to the effect that John is happy amounts to ascribing to the believer the property:

Being such that John is happy.

This treatment furnishes the basis for a well-known criticism by Stalnaker (1981), who considers that the property theory can only account for the transmission of beliefs indirectly, through inference. The property theory of belief started with Lewis's (1979) well-known Two Gods argument is here briefly summarized. We have the case of two gods, who inhabit a certain possible world. They know all the propositions true at their world. So they are omniscient. Still they suffer from ignorance: neither one knows which of the two he is. One lives on the tallest mountain; the other on the coldest mountain. Neither one knows where he lives. We have got the following inferential steps:

1. Each of the two gods believes every true proposition, but could have a true belief that he does not actually have;
2. If (1), then there can be beliefs the contents of which are not propositions;
3. If there can be beliefs the contents of which are not propositions, then the property theory is true;
4. The property theory is true (from 1-3). (See Feit 2008, 35)

The most commonly accepted forms of the property theory are based on Chisholm's idea that to have a belief x must bear a unique relation R to an object y and x self-

ascribes the property bearing R uniquely to something that has F. Thus we can analyse ‘John believes that Ann is clever’ as ‘John bears R uniquely to Ann and he attributes to himself the property of bearing R uniquely to something that is clever. (See Feit (2008) for developments of this idea).

The most crucial objections to this property-theoretic treatment are two. The first one is that in transmitting a thought, one must resort to inference. Thus the thought ‘You are clever’ is transmitted, say, by an inference from ‘the speaker is addressing a person y to whom he bears R uniquely and ascribes himself the property of addressing a y that has property F’ to ‘I am being addressed by a speaker who bears R uniquely to me and ascribes himself the property of addressing a y that has property F’, therefore I have F. The second objection has to do with verbs such as ‘want’, which are also considered verbs of propositional attitude. I may, for example, want to cease to exist and, nevertheless want all good people to be happy. On the property view, my second want has to be expressed as ‘I want to be such that all good people are happy and I want to be such that I do not exist’ – contradictory wants, it appears. (See Feit 2008 for discussion of these problems. For further discussion of objections to the property theory see Tomberlin (1991)).

It is not possible to review in detail all work on ‘de se’ attitudes – there is, for example the important work by Corazza (2004) on essential indexicals and logophoricity, as well Safir’s further work on logophoricity (Safir 2004). We cannot do justice in a brief chapter to all the work undertaken on the subject.

1.9 Philosophical ramifications of the ‘de se’ issue.

There is a growing literature on ‘de se’ attitudes. It is impossible to summarize it all. Suffice it to say that Stalnaker has recently used the debate on ‘de se’ attitudes to resolve Jackson’s problem based on a woman who has theoretical knowledge of certain aspects of the world, but no experience of it. Despite the fact that this is a brilliant philosophical book, I cannot review it here, as it has greater connections with philosophical, rather than with my linguistic interests (see my review in Capone 2009b).

2. Linguistic considerations by Higginbotham (2003).

So far, we have only considered philosophical treatments of ‘de se’ attitudes. At this stage, I propose to discuss Higginbotham’s views, both because they link with Lewis and Shoemaker’s ideas, and because they provide an analysis which makes it particularly clear and vivid that a ‘de se’ attitude entails a ‘de re’ attitude, which is what we require for our analysis based on informativeness and pragmatic scales or on contextual effects and processing efforts. Higginbotham (2003) considers a range of data such as the following:

- (17) John expects to win
- (18) John expects that he will win;
- (19) John expects that he himself will win.

Higginbotham considers that (18) does not necessarily have a ‘de se’ interpretation, while (17) and (19) necessarily have a ‘de se’ interpretation. He also says that syntactic constructions with PRO (where PRO is anaphoric) are even more first-personal than constructions such as (19). There is an ambiguity about (18) that allows the possibility of a ‘de re’ interpretation as well (albeit the ‘de se’ interpretation is preferred, and this fact demands a pragmatic explanation). Higginbotham makes use of Peacocke’s (1981) idea of a ‘de se’ mode of presentation:

Suppose that there is a special mode of presentation ‘self’ that a thinking subject *x* can use in thinking of himself, but not in thinking of people other than himself, and that others cannot use in thinking of *x*. A ‘de se’ thought will use an occurrence of [self_{*x*}] indexed to *x*.

The constructions that host ‘de se’ modes of presentation include verbs such as ‘imagine’, ‘remember’, ‘dream’, ‘pretend’, ‘know oneself’, etc. Higginbotham compares the following sentence types:

- (20) John remembered [his going to the movie];
- (21) John remembered [him going to the movie];
- (22) John remembered [himself going to the movie];

(23) John remembered [PRO going to the movie].

Unlike the other cases, (22) and (23) report ‘de se’ thoughts.

Given these facts, Higginbotham shows that the validity of the following deductive argument crucially depends on the presence of PRO, if a pronominal were substituted for PRO, it would become invalid:

Only Churchill gave the speech.

Churchill remembers [PRO giving his speech]; therefore

...

Only Churchill remembers [PRO giving his speech].

If we replace ‘Only Churchill remembers giving his speech’ with ‘Only Churchill remembers his giving his speech’, the argument is not valid.

An important linguistic fact noted by Higginbotham is that gerundive complements of ‘remember’ are associated with particular interpretations, according to which the remembered event is a **perceived** event. Thus, there is a difference between

(24) I remember giving a lecture at the University of Messina on 3rd November 1988;

(25) I remember that I gave a lecture at the University of Messina on 3rd November 1988.

I remember the event of the lecture through my **direct experience** of the event, given the semantics of (24); instead, I may merely remember that the event as described in the complement of ‘remember’ in (25) through someone else’s assertion, given the semantic import of (25).

To corroborate the considerations above, Higginbotham uses the example below:

(26) My grandfather died before I was born. I remember that he was called ‘Rufus’. But I do not remember his being called ‘Rufus’.

If ‘remember that’ and ‘remember + gerund’ had the same semantic import (if they contributed in the same way to truth conditions), then (26) would have to be a logical contradiction. But it is not. Hence the two constructions are associated with different truth conditional import. Higginbotham draws our attention to the following minimal pair:

(27) I used to remember that I walked to school in the fifth grade, but I no longer remember it;

(28) I used to remember walking to school in the fifth grade, but I no longer remember it .

Unlike (28), (29) is acceptable for Higginbotham. (28) reminds us of Moore’s paradox (Of course, to see why there is a problem in (28) one needs to stress that ‘remember’ is factive and that the assertion amounts to something like ‘I walked to school in the fifth grade but I no longer remember it’. (29) is acceptable, provided that we enlarge the scenario to include someone who said ‘You used to remember walking to school in the fifth grade’. The speaker of (29) says that he no longer remember the event in question, while he implicitly attributes responsibility for the truth of his remembering the event in the past to someone else who can report such an event of remembering.

I was told this example (which, I want to stress, belongs to Higginbotham, not to me) has problems, since it is not acceptable. S/he says that, if the utterances is acceptable, then one tends to read it (in terms of its internal dimension) as a direct experience of someone the memory of which can fade away with time or because of his partial election. I quite agree that one can have doubts on the grammaticality of (28), and thus propose, to remedy the problem, to consider it a loose usage (see Burge 2003 on loose uses of ‘remember’). In any case, the possibility ‘I used to remember walking to school in the fifth grade but I no longer remember it well’ is perfectly grammatical. This usage points to the fact that the internal dimension of PRO can be more or less fine-grained, a point which will be of use when I specifically deal with the internal dimension of PRO in terms of pragmatics.

Another point Higginbotham makes is that ‘de se’ constructions seem to involve immunity to error through misidentification (see work by Shoemaker). In other words, a person who says (29)

(29) I remember walking in Oxford

may be wrong on the place of the walking but not on the fact that it is his own walking that he remembers (leaving aside quasi-memories, cases in which someone else’s memories are implanted in a person’s brain).

Let us now see how Higginbotham characterises ‘de se’ attitudes semantically. He does that by making use of theta-roles as well as the Davidsonian’s idea that verbs have a hidden argument for events in logical form. (I assume readers are familiar with such semantics). Basically, Higginbotham’s idea is to identify the external argument of the verb of the complement clause with the external argument of the verb of propositional attitude. So the idea is that if I remember walking in Oxford, the agent of the walking is identical with the agent of the remembering. There is no such identification if the construction does not express a ‘de se’ concept as in ‘John remembers that he walked in Oxford’.

Consider the two cases:

(30) John expects to win;

(31) John expects that he will win.

Since there is no identification between the external argument of ‘win’ and the external argument of ‘expect’ in (31), we will represent (31) as

(32) (For John = x) $(\exists e) \text{ expect } [x, e, \wedge (\exists e') \text{ win } (x, e')]$.

(The approach considers propositions as sets of possible worlds à la Stalnaker; \wedge signals intensional abstraction over possible worlds).

Instead, (30) will be represented as (33)

(33) (For John = x) $(\exists e) \text{ expect } [x, e, \wedge (\exists e') \text{ win } (\sigma(e)), e']]$.

(32) represents a Russellian proposition as embedded in the matrix verb; (33) represents a mode of presentation that is first-personal in the sense of Peacock (1981). Since the identification of thematic roles has implications for reference as well, the Russellian proposition of (32) is expressed as a logical implication of (33) (in other words we expect (33) to entail (32)).

According to Higginbotham PRO in control structures embedded in verbs such as ‘remember’ also signals an internal dimension. When I say that I remember that I fell downstairs, there is no implication that my memory comes from my experience as the person who undergoes the falling downstairs. Someone else may have told me that I fell downstairs. However, if I say that I remember falling downstairs, I imply that I experienced the event and that I was involved in the event say as patient, the person who is affected by the event (we skip the issue of quasi-memories). This is the internal dimension of the remembering – I remember the event from the inside, as the person who was affected by the event. (So if the event caused me some pain, I remember the pain. It is not like remembering the event through the external perception of the event, say in case it was possible to connect my perceptual system to a camera and annul all other perceptions. In case it was possible to annul all my perceptions except for the visual images coming from a connected camera, it would not be true to say that I remember falling downstairs, but one could report that by saying that I remember that I fell downstairs. I take up this point in a critical discussion later on).

In order to represent the internal dimension of PRO, Higginbotham represents (34) as (35).

(34) John remembers falling downstairs;

(35) For John = x) $(\exists e)$ remember $[x, e, \wedge (\exists e')$ fall downstairs $(\sigma(e) \ \& \ \theta(e'))$, $e']$.

In other words, the falling downstairs is remembered as an event undergone by the person who remembers as a thematic role affected by the event of falling downstairs.

Consider now the case of the mad Heimson who believes to be Hume (Lewis 1979). We wonder if Heimson and Hume numerically have the same beliefs.

Consider ‘Heimson believes that he is Hume’ and ‘Hume believes that he is Heimson’ according to Higginbotham.

(36) (For $x = \text{Heimson}$), $((\exists e) \text{ believes } [x, e, \wedge (\exists e') \text{ be identical } (\sigma(e) = \theta(e'))], \ll \text{Hume}, e']]$

(37) (For $x = \text{Hume}$, $((\exists e) \text{ believes } [x, e, \wedge (\exists e') \text{ be identical } (\sigma(e)) = \theta(e')], \text{Hume}, e']]$

According to such readings, Heimson and Hume do not have numerically the same beliefs (given the identification of the believer and the bearer of the internal perspective, it has to be excluded that Heimson can be both the believer and the bearer of the internal perspective of the person identical with Hume).

2.1 Against Higginbotham (2003).

Safir (2004) argues that PRO is not necessarily associated with ‘de se’ interpretations in gerundival complements of verbs like ‘remember’. The examples used by Safir are the following:

(38) Oedipus tried to PRO commit patricide (and succeeded)

(39) Oedipus wanted/hoped to commit incest.

Safir considers that from the point of view of the omniscient narrator it is correct to say (38), (39) even if they do not reflect the following situations

(40) Oedipus attempted “I will commit patricide’

(41) Oedipus has a desire/hope “I will commit incest”.

According to Safir, the examples (38), (39) invariably show that PRO is not necessarily semantically associated with first-personal interpretations. Safir proposes that the association between PRO and a ‘de se’ interpretation is the consequence of a relation of non distinction between the controller and the event described in the sentence embedded in the verb of propositional attitude.

I think that if Safir’s examples were good, we could find evidence in favour of a conversational implicature analysis (a standard implicature) given that in standard

cases the implicature gets through, but in some cases it is defeated (contextually cancelled). I will take up this issue in my section on conversational implicatures and ‘de se’ readings. For the time being, I should note that Higginbotham will simply reply that the uses above (38), (39) prove nothing, since they are etiolated uses, things people say, without however wanting to say what is literally said. (In other words, such uses must be ungrammatical according to Higginbotham). I do not dissent from Higginbotham’s view, although I should say that there are other examples similar to Safir’s that deserve discussion (the reader has to wait until the section on conversational implicatures and ‘de se’ readings).

3. Pragmatic intrusion into truth-conditional semantics.

Although various authors have written about the role played by pragmatic inference in constructing a propositional form (e.g. Bach (1994), Levinson (2000), Recanati (2004), Stainton (1994), Bezuidenhout (1997), Powell (2001)), in this chapter I shall concentrate on the position of Relevance Theory on the semantics/pragmatics debate (mainly Carston 2002 and Wilson & Sperber 2002). As Horn points out (2004, 18) “taking the lead from work by Atlas (1979), relevance theorists have argued that the pragmatic reasoning used to compute implicated meaning must also be invoked to flesh out underspecified propositions in which the semantic meaning contributed by the linguistic expression itself is insufficient to yield a proper accounting of truth-conditional content”. Carston’s and Wilson & Sperber’s idea of pragmatic contribution to the proposition expressed has something distinctive because, unlike Bach, they believe that pragmatics contributes to what is said and, unlike Levinson (2000), they believe that the inferences that develop logical forms into propositional forms are explicatures, not implicatures¹. Carston’s and Wilson & Sperber’s ideas are similar to Stainton’s and Recanati’s, but they differ as to detail. For a review of intrusionistic perspectives, see Capone (2006, 2009).

3.1 Pragmatics, relevance theory and modularity.

Gricean pragmatics has been based on the so called Grice’s maxims: Quantity (informativeness), Relation (relevance), Quality (truthfulness), and Manner.

According to Grice, the speakers have to provide the right amount of information (neither too much nor too little) (Quantity), the information should be relevant (Relation), and true (Quality), and the conversational contribution should be as brief and clear as possible (Manner).

The Gricean program has been inherited by two schools of thought: neo-Griceans and Relevance theorists. Although I will not say much about the neo-Gricean approach, at least the following must be said to pay respect to this important strand of research. Neo-Griceans (Horn 1972, 1989, 2006; Levinson 2000) have decided to divide up the Quantity maxim into two, one stating that the speaker should say as much as she can (the Q-Principle), the other (R for Horn, I for Levinson) stating that she should not state the obvious. Levinson adds the M-principle to this picture, saying that if a speaker uses a marked expression, he wants to deviate from the implicatures associated with the use of the unmarked expression that could have been used instead (it generates the complementary implicatures) (see also Huang 2007).

Relevance theory can be seen as an attempt to elaborate on Grice's central idea that an essential feature of most human communication is the expression and recognition of intentions (Sperber & Wilson 1986/1995; Wilson & Sperber 2004; Blakemore 2000). The goal of inferential pragmatics is to explain how the hearer infers the speaker's meaning on the basis of the linguistic evidence provided. An utterance provides a linguistically encoded piece of evidence and verbal comprehension, thus, involving an element of decoding. However, the decoded linguistic meaning is just one of the inputs to a inferential process that yields an interpretation of the speaker's meaning. Given that the elements of meaning encoded linguistically in the sentences used do not exhaust the meanings inferred, and given that communication does not proceed by telepathy, the task of pragmatics is to explain in detail how comprehension takes place despite the relatively thinness (or poverty) of the linguistic input.

Relevance theory is not just a linguistic theory, but it is a more general theory of cognition. As Wilson & Sperber (2004) say, "relevance is a potential property not only of utterance and other observable phenomena, but of thoughts, memories, and conclusions of inferences. According to relevance theory, any external stimulus or internal representation which provides an input to cognitive processes may be relevant to an individual at some time" (p. 608).

The basic following principles are central for Relevance Theory:

Cognitive principle of Relevance

Human cognition tends to be geared to the maximization of relevance.

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.

Presumption of optimal relevance

- a. The ostensive stimulus is relevant enough to be worth the auditor's processing effort.
- b. It is the most relevant one compatible with the communicator's abilities and preferences.

Relevance-theoretic comprehension procedure

- a. Follow a path of least effort in computing cognitive effects: Test interpretative hypotheses (disambiguations, reference resolutions, implicatures, etc.) in order of accessibility.
- b. Stop when your expectations of relevance are satisfied (or abandoned).

(From Wilson & Sperber 2004).

3.2 Relevance theory and the modularity of the mind.

Relevance theory adopts a theory of the mind based on massive modularity. The modules that are involved in pragmatics are of course:

The language module, presumably innate;

The lexicon, a module that is structured and saturated through learning (Carruthers 2006);

The mind-reading module and its sub-module consisting of fast-and-frugal heuristics such as the principle of relevance.

In a series of articles Sperber and Wilson have tied their theory to massive modularity (but see also Carston 2002, in her lengthy introduction) (Sperber 1995; Sperber & Wilson 2002; Wilson 2005) and have connected their principle of Relevance to fast-and-frugal heuristics (see Gigerenzer et al. 1999).

The most exhaustive exposition of massive modularity is the one by Carruthers (2006).² According to Carruthers, the mind consists of a number of modules, functional components which are autonomous as far as their processing mechanisms are concerned, and sufficiently isolable and separate to allow the whole mind to execute its work, even if one of its components (or modules) has been damaged (say, the case of a person affected by a stroke, who will lose some of the functions of her mind (e.g. writing abilities), but not all of them). Modularity is also considered in relation to evolution, and the modular mind is considered capable of allowing modifications of one of its modules, or of adding a module as a consequence of adaptive changes in response to environmental challenges.

In the picture of massive modularity presented by Carruthers, each module is in constant dialogue with the other modules, in the sense that it can query a number of modules for information. Each module is switched on by a limited type of input (so there are heavy restrictions to the type of input a module can process, but surely a module can receive its input from different modules, as in the case of the practical reasoning module, which takes as input both desires and beliefs). Carruthers rules out categorically that each module has access to the internal operations of the other modules, as this would involve a simulation process imposing an excessive burden on processing. However, he admits the possibility that each module can search the archives of other modules.

² This departs somewhat from Fodor (1983), according to whom modular processes are supposed to be domain-specific, encapsulated, fast, mandatory, inaccessible to consciousness, neutrally localizable, and functionally dissociable. The main difference lies in the less stringent application of the notion of encapsulation and domain-specificity by Carruthers, as well as in the proliferation of modules, and in the insistence that central intelligence is nothing but the integration of all the various modules, plus some specific predispositions.

The same module can be located in more than one area of the brain, but is strictly correlated with a certain function. If two functions are structurally different, in that they involve different types of operations, then they will correlate with distinct modules. A module may consist of a number of sub-modules. For example, the mind-reading module consists of a module that deals with automatic inferences, as the ones studied by relevance theorists, and a reasoning module that deals with non-automatic, conscious and relatively slow inferences.

Modules may share a number of parts, and it is predicted that modules situated in areas of the brain that are sufficiently close will share a greater number of parts than modules situated at a distance, since the relatively slow propagation of the information between modules may offset the advantage of having parts that are shared. Carruthers believes that the same packages may underpin distinct modules. (This may be a drawback. After all, the modular architecture has the advantage of insulating one region of the brain from another and of preventing damage in one area from affecting the normal functioning of the remaining modules. If modules share parts, then how should we explain the fact that damage in one module will not affect the functioning of another? By saying that only modules situated in areas of the brain that are in close contact share parts, Carruthers limits the problem, as damage in one area will have consequences ONLY on neighboring areas).

Carruthers claims that central intelligence is the result of the interaction among modules, as integrated by further predispositions. Carruthers believes that central intelligence consists of two systems. In system 1, operations are highly specialized, and this is the reason why its output is fast. System 2 is responsible for steps in reasoning that are conscious. And that is what the global architecture does for us when supplemented with mental rehearsal and language.

Carruthers believes that there is a module that integrates information coming from distinct modules. He does not believe that this is the mind-reading module, but thinks that this is the language module. In fact, the language module is already capable of interfacing with modules using distinct algorithms. This is a reasonable hypothesis. Notice what happens in an utterance of 'Mary is drawing a triangle'. Here the speaker and the hearer are integrating information from the perceptual

module (the reference of ‘Mary’) and of the geometry module (a triangle activates geometrical information)³.

According to Carruthers, it is the syntax of language, as well as its recursive and hierarchical organization, that allows the integration of different types of information coming from distinct modules. For example, the fact that we can insert multiple modifiers under an NP node allows us to integrate information from different modules. NPs typically allow us to integrate information coming from perception (see also Jaszczolt 1999 on default semantics) with other type of information.

Another very important part of Carruthers’ research is constituted by his ideas on practical reasoning and mind-reading, namely that human beings, but also certain apes, are able to rehearse an action in their minds, broadcasting the result to central/conceptual modules, which give further beliefs and motivational states as output. A child who holds a banana in his hand may associate the banana with the concept of a telephone and then rehearse in imagination the action of phoning his grandmother, which provokes pleasure and thus determines the continuation of the action of rehearsing.

This book will have fruitful ramifications in due time and it will serve to throw further light on pragmatic theories of communication. Especially promising is the idea that a module – presumably the language module – is dedicated to integrating information coming from distinct modules. Seen in this light, an utterance can be seen as a merger representation à la Jaszczolt (2005), in which semantic, pragmatic and encyclopedic information merge through slots provided by the syntax present in the sentence or by syntactic constituents that are added at the level of the merger representation in order to make a thought plausible and maximally rational. Since perceptual, mind-reading, geometrical, etc. modules may be involved in communication of information, syntax is what allows the integration of the various types of information, and it is possible that syntactic constituents are added in the merger representation through the pragmatics of language in order to give room to constituents of thought such as modes of presentation of referents (see Jaszczolt 1999), thus integrating the theory of mind module with the perceptual module.

³ This example is my own, not Carruthers’s.

3.3 Relevance theory and the speaker's meaning

A number of authors, including Saul (2004) think that the exaggerated emphasis by relevance theorists on the hearer's interpretation process is a serious limit, as, after all, a pragmatic theory ought to be interested primarily in a theory of interpretation of the speaker's meaning. Now surely there is some truth in this criticism, and I must say that some relevance theorists such as Wedgwood (2007) insist that relevance theory is a theory about the hearer's interpretation of utterances. The question that is of real interest, in my view, is not whether RT is a theory of the hearer's interpretation, but whether it should be. My personal view is that relevance theory should be amended to be in line with more typically Gricean theories of the speaker's meaning. To neutralize possible distortions created by the emphasis on the hearer's processing of an utterance, I propose to follow Jaszczolt (1999) in requiring that interpretation follow the social, and not the individual, path of interpretation. To ensure that this obtains, we need some amendment of the principle of Relevance:

Presumption of optimal relevance

- a. The ostensive stimulus is relevant enough to be worth the auditor's processing effort.
- b. It is the most relevant one compatible with the communicator's abilities and preferences.
- c. It is the most relevant one compatible with the various clues and cues conducive to the speaker's intended meaning;
- d. It is the most relevant one compatible with the social path of intentionality.

I think Wedgwood's insistence on the hearer's perspective does more harm than good to RT. If I were challenged to provide arguments for this persuasion, I should quote from Capone (2009 forthcoming) on indirect reports. In this chapter I argue that the process of interpreting indirect reports usually (albeit not always, as interpretation process must adapt to special contexts) prevents the speaker (the reporter) from using an NP (coextensive with the NP actually used in the original utterance) that would transform (and distort) the original speech act, as such a move would place excessive processing efforts on the hearer. Now, my understanding of indirect reports is that pragmatic principles constrain both the speaker's behaviour and the hearer's behaviour. Emphasis on the hearer would surely provide only half of

the explanation required, perhaps the half which we need least, as after all it is legitimate to try and explain the speaker's behaviour as well, and such a behaviour clearly incorporates expectations about the hearer's behaviour.

Wilson & Sperber write:

The central claim of relevance theory is that the expectations of relevance raised by an utterance are precise and predictable enough to guide the hearer toward the speaker's meaning (p. 607).

Since an utterance creates expectations of relevance, it is licit to deduce that the speaker's linguistic behaviour is what creates such expectations. Thus, both he speaker and the hearer must be taken into account in an explanation of how relevance works. After all, when the authors say that human cognition tends to be geared to the maximization of relevance, 'maximization' is best understood as maximization by both the speaker and the hearer.

4. Beliefs 'de se' and pragmatic intrusion

In this section I consider 'de se' readings of attitude constructions and, in particular, constructions like 'John remembers walking in Oxford', 'John remembers he walked in Oxford', 'John remembers his walking in Oxford', 'John remembers he himself walking in Oxford'. My analysis starts with control structures like 'John remembers walking in Oxford' and then proceeds with the remaining constructions. Control structures in their 'de se' construals are determined through the semantics (I assume the truth of a story like Higginbotham's but I then slightly modify it). The remaining constructions are discussed in terms of pragmatics.

Since this is a rather complex and intricate section, we need sign-posts for readers here, to make sure that the analysis is taken for what it is, and not for what it is not. What I want to show in this section is that in some cases, but not in all cases, it is possible to derive the 'de se' interpretation through pragmatics

As far as constructions with PRO (such as 'John remembered going to the cinema') are concerned, I accept Higginbotham's story and claim that the external interpretation of PRO is semantic and first-personal. I want to distinguish however the concept of first-personal from the concept of using modes of presentation like 'I'.

A thought can be first-personal even if the speaker in talk with himself uses a mode of presentation like ‘You’, where by ‘You’ he, of course, means ‘I’. I present some arguments against the semantic analysis of the external interpretation of PRO, but conclude that these do not stand. PRO is first-personal. In particular, I use an argument by Feit (personal communication) to show that PRO must be first-personal and that Higginbotham’s semantic analysis probably needs further tightening up. I also use a circularity problem, to show that Higginbotham’s syntactic analysis is, after all, presupposed in my pragmatic analysis.

I argue that the cases where Castañeda used the asterisk are cases in which a pronominal, which is not PRO, are assigned a ‘de se’ interpretation (in due time I will also claim that this pragmatic process involves using (at the level of a merger representation a la Jaszczolt 2005) a PRO construction like the ones studied by Higginbotham in combination with verbs of imagination, remembering, expecting, etc.). A pragmatic story is used for sentences like ‘John remembered his going to the movie’ – with the difference that here I argue that ‘his’ is not assigned an asterisk à la Castañeda.

The internal dimension of PRO is a separate question from the external first-personal interpretation of PRO. While the implicature analysis has very limited effects on the external dimension of PRO, since I have accepted that PRO is first-personal through semantics, I argue that the internal dimension of PRO is conveyed not through semantics, but through implicature (or explicature).

The pragmatic analysis has the following structure:

- a. Analysis of inferential enrichments amounting to ‘de se’ interpretations;
- b. Analysis of pronominals used instead of PRO in control structures;
- c. Analysis of reflexives used in control structures;
- d. Apparently ‘de se’ uses of pronominals with attributive construals.

In each of this section I will substantiate the claim that ‘de se’ constructions are intrusive ones and that pragmatics serves to resolve interpretative ambiguity and to determine a full proposition.

4. 1 Mode of presentations of first-personal readings: semantics or pragmatics?

Before we proceed with our pragmatic story, it will be important to characterise ‘de se’ pragmatic interpretations more fully. What kind of representation must be part of the explicature when a ‘de se’ thought is involved? Presumably, when the speaker says:

(42) *Giovanni crede di essere felice* (John believes he is clever)

there is an inference that the speaker has the following mental representation:

“I am clever”.

This is on top of the semantics provided by Higginbotham for controlled clauses of attitude verbs⁴. In fact, it would be possible, strictly speaking, for the semantic interpretations by Higginbotham to be accessible to the believer without his using a mental occurrence of ‘I’. Higginbotham’s story could be true even if the believer thought of himself as the believer of his thought, without ever pronouncing (or using a mental occurrence) of the word ‘I’. However, it is reasonable to suppose, accepting the considerations by Millikan, that ‘de se’ readings involve a mode of presentation that somehow incorporates ‘I’. Sentences such as (43)

(43) *John thought he was clever*

are ‘de se’ in the sense of incorporating mental linguistic materials such as ‘I’ when it is clear in context that the evidence for the thought comes from the believer having uttered a statement about his/her belief.

But is there a sharp difference between this additional pragmatic component and Higginbotham’s semantics? It is true that Higginbotham does not explicitly consider utterances of ‘believe’ in connection with PRO, but since in Italian belief-constructions obligatorily involve PRO (in the ‘de se’ interpretation) we must assume an extension of Higginbotham’s story. Higginbotham’s extended treatment would have to amount to including a use of the ‘believer of his/her thought’ in a belief

⁴ One should note that Higginbotham does not extend his analysis of ‘X remembers walking in Oxford’ to beliefs, but the extension is required for languages like Italian, which, unlike English, has control structures embedded by ‘believe’.

attribution (e.g. John believed that the believer of his thought was happy) – strictly speaking it involves usage of temporal variables as in (44)

(44) Giovanni credeva di essere felice (John believed he was happy)
John believed at t that the believer of this thought at t was happy,⁵

which presupposes that if John believes at t thought x , he cannot believe at t a thought y , y distinct from x . (Can one have two distinct thoughts at the same time? Presumably not).

My own addition to his treatment says that on top of Higginbotham's semantics, there is an inference to the effect that the believer makes use of a mental occurrence of the word 'I' – he says 'I' in his mental sentences (provided that the context is the right one). Now, if the mental occurrence of 'I' were identical with Higginbotham's contrived solution 'The believer of his/her thought at t ', obviously there would be no reason for having this additional pragmatic component.

A cogent reason for opting for my own treatment is given by Feit (2009, personal communication):

Another reason that I do not think Higginbotham's account can handle 'de se' cases adequately is this. It seems possible that somebody could believe (correctly or mistakenly, it does not matter) that he is not the only thinker of a certain thought, for example he might believe that God is thinking it too. More generally, he might think that he is not the only thinker of any of his thoughts. But, even with this, it seems he could have a 'de se' belief. But on Higginbotham's view – and other similar views – such a belief amounts to "the believer of this thought is F." This cannot be what the belief amounts to however, since he does not think there is a unique believer, the believer, of his thought. Moreover, if someone else (God perhaps) really is having the same thoughts, then all Higginbotham-style beliefs are false, but he could surely have some true de se beliefs nonetheless.

⁵ It might be said that Higginbotham does not particularly discuss this example. Yet, it is natural to think that he must accept this semantic analysis of the Italian example because of his commitments concerning 'John remembers going to the cinema'.

Thus a minimal requirement for making sense of ‘de se’ attitudes is to say that the mental occurrence of ‘I’ (say in mentalese, see Feit 2008 on this) must be a demonstrative along the lines of Gareth Evans (1982) and Perconti (2008). As Evans says, the demonstrative identification does not go through the recognition of any property⁶. But is there something that the word ‘I’ can refer to? Evans argues that there is substantive content to our ‘I’-ideas. While for philosophers such as Strawson to the judgement that I am in pain there does not appear to be anything corresponding to the identification of something that is in pain (the judgement might as well be expressed by ‘There is pain’) , for Evans by using ‘I’ we must identify with an element of the objective order. (See also Grush 2002, for the exegesis of Gareth Evans).

Now suppose Higginbotham replies:

All you have shown is that the first-personal interpretation of PRO needs to be grafted to the semantics I proposed, and one obvious way to do this is by placing in the semantics the further constraint that the mode of presentation of the agent of the believing or remembering a certain thought is in the first person. (This move is reminiscent of Harcourt and Millikan)

Presumably, we have to modify Higginbotham’s elucidations for an utterance of, say, ‘John believes he fell downstairs’ in the following way:

(45) For John = x) ($\exists e$) believes [x, e, \wedge ($\exists e'$) fall downstairs ($\sigma(e) \& \theta(e')$ & the mode of presentation of $\sigma(e) = 'I'$), e']].

After all, this is still a completely semantic meaning elucidation.

I suspect that such a reply would lead us further in the pragmatic direction. After all, suppose that John thinks the following thought: ‘Why don’t you go to the cinema?’⁷

⁶ See p. 170-1, Evans (1982) on demonstrative identification based on an information-link between the subject and object as well as on the ability to locate the object relative to egocentric space and to objective space.

⁷ Well, of course we do not ordinarily think this way. But it is not impossible to think this way. I have before my mind the situation of a person who slaps his face

while speaking to himself. Surely, it would not be misleading to report such a thought saying: ‘John considers going to the cinema’.

This is a first-personal thought, but no first-personal mode of presentation was used in thought.

Or suppose the thinker says to himself ‘Oh gosh, you are really smart!’. It would not be misleading, in this case, to report the thought: ‘John believes that he is really smart’, even if he did not use in his thought the mode of presentation ‘I’ (instead he used ‘You’).

Or suppose John says to himself ‘Smart, this criminal’, this being a reflection on himself, the rapist whom everybody in his town fears. Surely, it would not be false to say that John believes that he is really smart.

So, it is not necessary to use ‘I’ in thought to express a ‘de se’ thought (these are all cases in which in Italian one would use PRO instead of the bare pronominal in English). Since the inference that the mode of presentation of the believer in a ‘de se’ thought is ‘I’ is quite standard, but it can be cancelled in some cases, it is not unreasonable to say that the inference is pragmatic.

But this is not like saying that the interpretation of PRO in ‘de se’ constructions like those utilized by Higginbotham is not first-personal. We have to dissociate the concept first-personal from the actual usage of a certain first-person pronominal like ‘I’. An interpretation can be first-personal even if the speaker thinks of himself as ‘You’ (say when talking with oneself) or through an epithet (‘Bastard’). It is first-personal if the pronominal used is a device used by grammar for integrating information from the module Theory of Mind, and where the mind in question is the thinker’s. A compromise between my view and that of Higginbotham can be made by using the modularity of mind story I presented before. PRO works as a syntactic constituent that allows the integration of information coming from the theory of mind module.

Summing up, the sentence ‘John believes he fell down’ needs to be represented as the following:

saying ‘*Sei proprio stupido*’ (You are really an idiot). This is probably akin to the situations Goffman describe in ‘Forms of Talk’, in which one dissociates from his previous conduct (and from his previous self). These uses, anyway, are also attested in the literature on the modularity of the mind, albeit circumscribed to pathological situations (see Robbins 2007, 310).

For John = x) ($\exists e$) believes [x, e, \wedge ($\exists e'$) fall downstairs ($\sigma(e)$ & $\theta(e')$) & the mode of presentation of $\sigma(e)$ = any pronominal or mode of presentation that is a suitable transformation of 'I' and ultimately reducible to 'I' , e')].

One further reason for adhering to a semantic story of PRO and for not wanting to say that PRO conversationally implicates or is associated via an explicature to its first-personal reading is that the explicature triggered by example (50), as I will claim later, uses syntactic information, and, in particular, the possibility of expressing the 'de se' reading through PRO. If we say that the 'de se' reading of PRO is conversationally implicated (or alternatively explicated), then we are at a loss when we want to explain the 'de se' reading of (50) through the syntax of e.g. 'Giovanni crede di essere Italiano'. We would have to say that 'de se' concepts are completely pragmatic, but such a story would have no points of contact with the views on the modularity of mind. It is, on the contrary, plausible that a theory of mind module is at work in 'de se' readings and that this is the reason why 'de se' interpretations should be correlated with a special syntactic construction.

4.2 Towards pragmatics: Castañeda's example.

Let us now consider Castañeda's influential example:

(46) The editor of Soul believes that he* is a millionaire

Unlike the philosopher's language, ordinary language has no asterisks. I agree that the preferred interpretation is one according to which the editor of Soul believes that he himself is a millionaire, but this is not a matter of semantics, as there is an alternative reading according to which the interpretation is not 'de se'. Suppose, for example, that the editor of Soul believes of the person (himself) he sees in the mirror that he is a millionaire (while, for some reason he does not recognize his familiar face). A sentence such as (47)

(47) The editor of Soul believes that he is a millionaire

is suited to expressing the speaker's meaning – however, no 'de se' reading is intended in this case (We agree, the example is contrived, and is based on philosophical sophistication, however it is not an impossibility). The interpretation in (47) where an asterisk is used to signal pragmatic disambiguation must not be taken for granted, but is the result of cognitive processes at work.

We may also want to explain Perry's example:

(48) The dean was surprised to find that he believed himself to be overpaid.

In a situation in which the dean believes that all professors who publish less than ten papers per year are overpaid (but forgets that he himself has published less than ten papers), a speaker may utter (48). Linguists may have reservations about such an example. They may feel it is contrived or that this is a loose or etiolated language use⁸. Whether the use is correct or illegitimate (strictly speaking), we have to explain such a use as well by resorting to a pragmatic theory. While in the case of (47) we must explain why a 'de se' reading accrues to the utterance, in the case of (48) we have to explain why a sentence/utterance which is typically associated with a 'de se' reading is divested from its ordinary interpretation. Obviously, while the pragmatic process at play in (47) is a case of a standard conversational implicature, the process involved in (48) is a case of a particularized implicature. The implicature overrides the usual semantic interpretation associated with the sentence ('de se' reading) (see Dascal 2003 on the divergence between sentence and speaker's meaning). I assume that the world knowledge against which the utterance of (51) is processed promotes the non-first personal reading. Given that we assume that the Dean thinks highly of himself and would never say of himself that he thinks he is overpaid, we assume that the interpretation of (48) is not a 'de se' one. The utterance comes to be interpreted as ironic, because, on the one hand, the speaker says that the dean believes that he himself is overpaid, on the other hand we know that the Dean would never think that of himself. The utterance is 'echoic' in that pragmatic interpretation construes it as what the Dean would have to say of himself if he were to accept what the other people believe of him. The 'de se' reading is a reading which expresses what the

⁸ Recanati (2007, 173) also believes that 'himself' is less first-personal than PRO. His example is: "John imagines himself being elected". Presumably (I infer this from the passage in Recanati's text), someone could say this without attributing a 'de se' attitude to John. There is no explanation about why this should be the case, though.

Dean would think of himself in a possible world in which he conforms to what other people think of him.

Anyway, I should say I am puzzled a bit by Perry's example. I think that what he wants to say requires a different example, such as 'The dean would have been surprised to find out that he believed himself to be overpaid'.

I guess that what Perry wanted to prove with this example is that 'de se' readings are in all cases pragmatic and not semantic. Could there be a pragmatic interpretation that is not founded on a semantic concept? In theory it is possible – as Recanati (2004), Carston (2002) and Wilson & Sperber say in their articles and books - that pragmatics furnishes new concepts on the basis of existing ones (a phenomenon called 'modulation'). So we should not discard a priori the possibility that 'de se' readings are only pragmatic interpretations, which make their way into language through grammaticalization (see Ariel 2008). However, it cannot be excluded that what started as pragmatics ended up as semantics or grammar (Levinson 2000). We shall explore possibilities open-minded.

How should a relevance-theoretic treatment of (47) proceed? I assume that the interpretation according to which the speaker attributes a belief 'de se' to the subject (of belief) is more informative than the 'de re' interpretation. We can reasonably assume that an interpretation that excludes a greater number of states of the world (see also Levinson 2000; see also Huang 2007) is more informative. It is reasonable to think that on a relevance-theoretic treatment this is true as well. What is to provide information, in fact? To provide information is to provide input to inferential processes, among which there is the strengthening of existing assumptions or the elimination of current assumptions or the creation of cognitive effects that would not derive from existing assumptions alone. A proposition that eliminates a greater number of states of the world is, ipso facto, more informative than a proposition that eliminates a fewer number of states of the world, because it either eliminates existing assumptions or interacts with existing assumptions in such a way as to provide a greater number of cognitive effects than the ones that would derive from existing assumptions alone. Suppose one knows that all students have arrived, rather than that some students have arrived. Suppose one furthermore knows that all students who have arrived will receive a present. Then one knows more than if one knows that some students have arrived. If all students consist of A, B, C, D, E, one derives greater cognitive effects from knowledge that all students arrived, since one

will be able to say that all of A, B, C, D, E will receive a present. Instead, having only knowledge that some students have arrived, it will not be possible to say which of A, B, C, D, E will receive a present.

Now let us go back to our 'de se' interpretation in (50). We have to ask ourselves which is more informative: the 'de se' or the 'de re' interpretation? Matters of entailment may decide the issue. Consider again Higginbotham's analysis of the 'de se' reading and of the non-de se reading:

(49) (For John = x) ($\exists e$) expect [$x, e, \wedge (\exists e')$ win ($\sigma(e)$), e']].

(49) represents the 'de se' reading of 'John expects to win'.

(50) (For John = x) ($\exists e$) expect [$x, e, \wedge (\exists e')$ win (x, e')].

(50) represents the non 'de se' reading (that is, the 'de re' reading) of 'John expects that he will win'. The person who is committed to the logical form (50) is committed to logical form (49), but it is not true the other way round. This means that the 'de se' reading entails the 'de re' reading. Since 'de se' readings entail 'de re' readings they are more informative.

We do not need to go through the entailment (or deduction) step to argue that the 'de se' reading is promoted by pragmatics to default interpretation, though. All we need to show is that the 'de se' reading has greater cognitive effects than the 'de re' reading, processing efforts being the same. I think we can concoct a philosophical story. Suppose that Mary believes she has to take a tablet at 9 in the morning (the usual tablet she takes daily). Suppose that such a tablet has an undesired effect m , which can be eliminated by taking tablet b (the same person who takes tablet a must take tablet b to avoid an unwanted effect m): Then the 'de se' reading of the sentence 'Maria believes she must take the tablet b ' has greater cognitive effects, since only in case Maria thinks of herself as herself she is interested in preventing the consequences of taking tablet a . Since the 'de se' reading has greater cognitive effects than the 'de re' one, which offset the processing costs incurred, it will be promoted by the Principle of Relevance.

4.3 The pragmatic process in belief sentences involves use of a PRO structure in the language of thought.

At this point I would like to make a bolder and more controversial claim. This claim comes from studying the Italian translations of the sentence (51):

(51 a) Il direttore di Soul crede che (egli) è un milionario (The editor of soul believes he is a millionaire);

(51 b) Il direttore di Soul crede di essere un milionario (literal: The editor of Soul believes PRO he be a millionaire).

It is interesting that in Italian the ‘de se’ reading cannot be expressed by (51 a), which is definitely suitable for the ‘de re’ reading. (Thanks to Daniele Gambarara for being adamant about this). Italian uses obligatorily (51 b) for the ‘de se’ reading. Why is it so? Why is it that Italian is so different from English in this respect. But is it so different? Suppose we accept Higginbotham’s story that PRO in ‘de se’ constructions is more first-personal than both ‘he’ and ‘he himself’. An immediate question arises. Why is this the case? A natural explanation comes from accepting Carruthers’s story that language serves as an integrator of modular information (Carruthers 2006). The recursiveness of syntax is taken by Carruthers as allowing the integration between inputs from distinct modules. A device for integration is taken by Carruthers to be the possibility of adding multiple modifiers within an NP mode. That some inter-modular connection, or interfacing is required, is also independently supported by Robbins (2007) in his important reply to Emma Borg. What matters for the purpose of this chapter is that Robbins claims that “there is nothing in the modularity story that rules out process-constitutive (or process-dependent) information flow between modules. The idea of intermodular connection is also supported by de Villiers (2007), who, citing St. Augustine, argues that information from the domain of Theory of Mind seeps into the process of simple word learning, given that “the human context of shared eye gaze and pointing helps delimit the possible meanings of a new word” (de Villiers 2007, 1864).

Suppose we take Carruthers’ story one step further and consider PRO (in subject control) as a null empty category that allows the integration of information coming

from the theory of mind module and from the perception module (we must not forget that even PRO must be coindexed with a subject). Then it follows that PRO is different from full pronominals, or from full NPs. Full NPs are referential, by default, as Jaszczolt (1999, 2005) says (see also Delfitto 2004). The reason why they are referential by default is not a story about informativity, but a story about modularity, as full NPs and pronominals allow the integration of information coming from the perception module, and it follows that information coming from the perception module must be referential by default. Now a little modular story is in place, which explains why English is not so different from Italian, after all. In fact, the interpretation of (50) intended as ‘de se’ (as in (50 b)), contains PRO at logical form in the merger representation (Jaszczolt 2005). Following Davidson and Higginbotham, we are prepared to accept that at logical form there may be implicit elements, such as, for example hidden quantifiers and positions for events as in

(52) John fell

There is an event **e**, such that John fell in **e**.

However, I depart from Higginbotham and Davidson by saying that in the cases I am discussing, the logical form that departs from the sentential structure ‘John believes he is clever’ is a constituent of what Jaszczolt calls the ‘merger representation’. By analogy, I propose that the ‘de se’ interpretation of (50) has a PRO at LF, and, thus, is not much different from Italian (50b) (Il direttore di Soul crede di essere un miliardario). I propose that the de se reading of (50) should always be represented as (53) in the logical form of the merger representation.

(53) The editor of Soul believes PRO be a millionaire.

I furthermore propose that the syntactic representation involved in this interpretative reading is articulated at the level of the merger representation, following K. Jaszczolt’s (2005) idea that compositionality is a characteristic of the merger representation. The merger representation of (53) contains the syntactic structure [PRO be a millionaire] and thus it is not necessary to strain our intuitions about the semantics of (53), as what is needed to obtain the right interpretation takes place at the level of the merger representation, including articulating the syntactic

representation containing PRO. Of course, we can retain the intuition that, after all, (53) must be considered as expressing two completely different sentences; however shifting the burden of explanation to the level of the merger representation allows us to preserve our intuitions about the surface syntax of English which contains no representations such as ‘John believes PRO be a millionaire’. (Instead notice that the sentence ‘There is an event e such that John fell in e ’ is fine in English and can be said to be part of the sentential logical form). What is not licit in the surface syntax is licit at the level of the merger representation, as we can assume that the merger representation can use a language of thought, which is universal, and which allows the translation of Italian sentences into English and vice versa. Suppose that this language of thought requires us to use PRO whenever a first-personal ‘de se’ reading occurs. Such a language of thought coincides with the surface structure of Italian, but is also what allows the translatability of an English ‘de se’ sentence into the Italian ‘de se’ sentence despite the difference in surface structure. This language of thought ensures that a language like English which has PRO in ‘de se’ constructions in the case of ‘remember’, ‘expect’, ‘imagine’ can utilize this construction in the language of thought in order to express the ‘de se’ interpretation of a sentence which does not utilize PRO at surface structure.

The story above encounters a severe objection. Consider Stanley’s (2007) First assumption:

First assumption:

In semantic interpretation, one may never postulate hidden structure that is inconsistent with correct syntactic theory (Stanley 2007, 35).

My treatment seems to run afoul from Stanley’s certainly reasonable assumption. After all, if constituents appear at the level of what Jaszczolt calls the ‘merger representation’ it is not unreasonable that they should be inserted in nodes that belong to syntactic trees and they should themselves have a syntactic articulation.

How can we reply to this objection?

There are two possibilities.

One is to say: Well, the language of thought need not follow the same rules of grammar as, say, English. I can very well think by using ungrammatical sentences.

The other is to adopt the considerations which I presented in Capone (2008), where I claimed, following Kent Bach, that we can explain opacity in belief sentences by using appositives. However, unlike Bach, I did not posit appositives under NP nodes, but I postulated a sentential appositive having an internal articulation and allowing places for NPs that were associated with modes of presentations, these modes of presentations being linked though pragmatics with the NPs in the embedded clause of a belief sentence. Now, in adopting that idea in this chapter too, I am not postulating anything ‘ad hoc’, for the purpose of rescuing the analysis above, interesting and problematic thought it is. The treatment is quite general.

Now suppose that in a sentence such as ‘She believes she is clever’ we do not insert at random the constituent PRO or, in any case, we do not replace the sentence at the level of the merger representation with the ungrammatical ‘Mary believes PRO to be clever’, but we posit structure in exactly the way I proposed in Capone (2008), but at a more abstract level. So we would have the following structure:

Mary believes [that [she is clever] [PRO be clever]]

This could not be analogous to: Mary is ready [to collect the insurance] and thus a legitimate question arises as to what is it that selects the [PRO be clever] clause. However, following de Vries (2002), we know that even sentences can have appositives, and this suffices to allow us to have the following structure:

[Mary believes she is clever] [Mary Attitude verb PRO be clever]

The only objection I can envision to the structure above is that we have a sentence with lexical NPs and VPs and an appositive having an abstract verb: however, since it is well accepted in semantics that there is a class of verbs called verbs of propositional attitude, that introduce opacity, and have well known semantic characteristics, it is not impossible for the human mind to replace a verb with a categorial type. In any case, Stanley’s constraint is only at the level of syntax, and he

does not say anything about having abstract categories at logical form (in any case, the fact that grammarians are at ease using abstract verbal categories is a proof that the mind is perfectly capable of creating an abstract category and use it in an appositive).

Now, what is the relationship between PRO and ‘she’? It is not a relationship of reference, but mainly a relationship such that the mode of presentation ‘PRO’ is transmitted from the appositive to ‘she’, through some kind of linking.

4.4 DE re interpretations: the pragmatic interpretations of pronominals, as used instead of PRO.

Let us see what happens if a full pronominal is used instead of PRO. So consider again the minimal pair from Higginbotham (2003):

(54) John remembered [his going to the movie];

(55) John remembered [PRO going to the movie].

Higginbotham says that PRO is associated with a ‘de se’ interpretation, while (54) is not. We ask why it should be the case that ‘John believes he* is clever’ is typically associated with a ‘de se’ reading, while (55) is not. Neo-Griceans (e.g. Huang 2000, 2007; Levinson 2000) have an easy explanation at hand. Suppose that

<his, PRO> form a Horn-scale, such that the two forms are from the same semantic field. Since PRO is associated with the ‘de se’ reading, it is more informative than the ‘de re’ reading. Thus PRO ends up entailing ‘his’. Use of ‘his’ at this point will implicate that the ‘de se’ interpretation does not obtain. (The only problem for this analysis is the equal lexicalization constraint: should we say that PRO and ‘his’ are equally lexicalized? This is a problematic choice).

Alternatively, one can say that the full pronominal is more marked than PRO and therefore triggers an M-implicature to the effect that the interpretation complementary to that of PRO takes place. (Remember, the M-Principle says that the usage of a marked expression instead of an unmarked one will trigger a

complementary implicature: the problem here is that PRO, if what Higginbotham says about the first-personal reading is correct, is not coextensive with the full pronominal – as required by the M-Principle (see Levinson 2000; Huang 2007, and references therein).

Both routes are not devoid of problems that need to be addressed somehow. Now, we want to find a plausible alternative relevance theory story. Suppose we say that the overt pronominal requires greater processing efforts than ‘PRO’. Then we require additional contextual effects to counterbalance the gratuitous processing efforts. There will be compensatory contextual effects if the interpretation is complementary to that of PRO (or even if it is distinct from that of PRO. Thus, the ‘de re’ interpretation, which is complementary to the ‘de se’ one, gets through.

4.5 The internal dimension of PRO: ‘remember’ and other verbs.

In this section I shall discuss verbs such as ‘remember’, ‘imagine’, ‘expect’, ‘dream’, ‘forget’ in terms of the internal dimension of PRO. Since such considerations are sparked by Higginbotham’s reflections on the internal dimension of PRO in connection with ‘remember’, I will start with ‘remember’, on which there is a wealth of philosophical considerations.

In particular, I will discuss what Higginbotham (2003), based on ideas by Martin & Deutscher (1966) and Shoemaker (1970) calls ‘remembering from the inside’ associating it with control structures (‘John remembers falling down the stairs). Following Norman Malcom (1963), Shoemaker distinguishes between the semantics of ‘John remembers that Caesar invaded Britain’ (factual memories) and ‘John remembers falling down the stairs’, the latter sentence being associated with something which one remembers happen, as a result of observation or experience.

Shoemaker (1970) discusses only cases like ‘John remembers walking in Oxford’ and says:

It is a necessary condition of its being true that a person remembers a given past event that he, the same person, should have observed or experienced the event, or known of it in some other direct way, at the time of its occurrence. I shall refer to this as the ‘previous awareness condition’ for remembering” (Shoemaker 1970, p. 269).

He adds that “When a person remembers a past event there is a correspondence between his present cognitive state and some past cognitive and sensory state of his that existed at the time of the remembered event and consisted in his experiencing the event or otherwise being aware of its occurrence (p. 271).

I take that the awareness condition and the correspondence condition for Shoemaker are entailed semantically by a sentence such as ‘John remembers falling down the stairs’ and they correspond more or less to what Higginbotham calls the internal dimension of PRO. Now, while my aim in this section is to argue that the internal dimension of PRO is conversationally implicated by sentences such as ‘John remembers falling down the stairs’, I want to do justice to the importance of Shoemaker’s considerations and suggest that the internal dimension of PRO may be more or less fine-grained and that conversational implicatures may be responsible for the more fine-grained dimension of the internal dimension, while we can assign semantics the task of doing justice to the considerations by Shoemaker, which seem to me to be not implausible. In particular, I think we can accept that in uttering a sentence such as (56):

(56) John remembers falling down the stairs

the awareness condition needs to be satisfied, and John’s memory needs to be caused by a perception of his experience of falling down. Furthermore, the correspondence condition, whereby there must be a correspondence between the memory and the experience or sensory state that existed at the time of the event, must be satisfied. If John remembers falling down, then there must be an experience which has triggered his memory – there is a rough correspondence between the experience and the memory. However, how fine-grained the correspondence ought to be has not been specified by Shoemaker. Is it not possible that only part of the experience has been recalled, thus making it possible that there is a correspondence between the sensory state of the event and the memory, even if we can admit that the fully articulated dimension of the sensory state has been communicated in a more fine-grained way through pragmatics?

I think it is not unreasonable to propose that the full internal dimension of PRO is communicated via pragmatic intrusion. When we say (57)

(57) John remembers falling downstairs

we surely mean that the John is remembering the event from the inside, that he was at the same time the perceiver of the event and the participant affected by the event (he did not just see the blood on his face, but he also felt the pain and the event of remembering the pain could occur only through the experience of the pain (his feeling his pain)). However, it is not necessary to place all burden of both first-personal and the internal dimension of PRO on semantics. The burden can be divided between semantics and pragmatics. After all, it is not unnatural to say:

(58) John remembers falling downstairs, but he does in an incomplete way. He does not remember the pain he felt. The memory is to him like a film he is watching.

Notice, in this case, both the awareness and the correspondence conditions proposed by Shoemaker are satisfied, even if some fine-grained aspects of the internal dimension have been lost.

The statement (58) could be justified, in case John has partial amnesia or has (voluntarily or involuntarily) erased parts of his experience, namely those memories which are more painful.

Also we can think of the case in which a memory is so painful that, although the person in question does remember the event (say, an accident), s/he does not want to recall it. By failing to recall its most painful parts, the memory will be partial.

After all it is not so unreasonable to assume that memories can fade away and that parts of them can be erased. So, the idea that the internal dimension of a memory can be erased (removed) is not so outlandish. Psychologists often say that women who gave birth to a child remove the pain from their memories – this is why they are willing to give birth to a second child, without much thought about it. Furthermore, going back to Higginbotham's example, partially modified:

(61) I used to remember walking to school in the fifth grade, but I no longer remember it very well

this example can also be understood as saying that the speaker had an exhaustive memory of the event of walking to school in the fifth grade, but now he no longer has this exhaustive memory (he has a partial memory)⁹. Memories can be partial – as parts of memories can be removed. However, in the typical case, the internal dimension of the memory does not disappear. So, if a person says ‘I remember falling downstairs’ the full internal dimension is communicated as well, but by pragmatics. Through a pragmatic increment, we build up the explicature.

Let us consider how Relevance Theory can deal with similar examples. Consider (59):

(59) I remember falling downstairs.

If one falls downstairs, in the prototypical case, one feels pain. It is, therefore, reasonable to assume that in addition to visual memories, the speaker has other types of memories: tactile memories for example (scratches, pain on his bones, etc). It is, therefore, probable that he is remembering the event from the inside. However, it is not implausible that the ‘internal dimension’ can be (partially) cancelled, probably because the memory can be as painful as the real experience which one is re-living. To put things in the words of Carruthers (2006), when one remembers an event, one rehearses the event in mind, thus evoking motor-sensory schemata that are broadcast to central/conceptual modules and may generate real pain and frustration.

At this point, let us follow Cimatti (2008), in the idea that the subject is constituted through speech and let us make use of the psychotherapy situation as a hypothetical situation. Suppose that the speaker says (62) in the course of a psychotherapy. The patient, who was pushed down the stairs by his mother, removes all sensations of pain. The aim of the therapy is to help the patient relive the situation and recuperate the important parts of the memory he has removed – say what modifying slightly Higginbotham’s terminology we could call the full internal

⁹ The example has the other reading noted by Higginbotham, as well.

dimension of the memory. Then this is a case in which the internal dimension of PRO has been partially suppressed and one tries to recuperate it. Thus, at the end of the psychotherapy the same sentence can be uttered having a different meaning including the full internal dimension as well. If the same sentence can be uttered at different moments by the same patient, thus rehearsing an experience and broadcasting motor/sensory schemata to the central/ conceptual systems, and can broadcast different schemata at different moments, thus provoking different corporeal sensations, this can be taken as proof that the **full** (or fully articulated) internal dimension of PRO is not associated with sentence meaning, but, at most, with utterance meaning, and, in particular, with the speaker's meaning.

The little – not too implausible – story above can show that the full internal dimension of PRO has greater contextual effects. By recuperating the internal dimension of PRO, the speaker can recuperate feelings that have consequences on parts of his personality. Alternatively, he can recuperate beliefs that, in conjunction with other beliefs, can produce further beliefs.

See in fact the following deduction:

John remembers falling downstairs.

If John remembers the event (fully) from the inside, he remembers feeling pain.

John remembers the event (fully) from the inside (premise furnished through pragmatics)

If he felt pain, he hated his mother who pushed him.

....

John hated his mother

Since the premises added by pragmatic inferences (in particular, the internal dimension of the memory) lead to further contextual effects through deduction, it is reasonable to accept that such inferences are motivated by the desire to be relevant, to create abundant cognitive effects with minimal cognitive processes¹⁰. As Wilson & Sperber (2004) say:

¹⁰ On unexpressed premises in enthymemes see Piazza (1995).

The most important type of contextual effect is a CONTEXTUAL IMPLICATION, a conclusion deducible from input and context together, but from neither input nor context alone. For example, on seeing my train arriving, I might look at my watch, access my knowledge of the train timetable, and derive the contextual implication that my train is late (...). (p. 608).

This topic seems to have intrigued an influential linguist like John Lyons, who notices a difference between

(60) I remembered closing the door

(61) I remembered myself closing the door. (These examples are numbered as (3) and (4) in Lyons' paper).

Lyons (1989) writes:

As to the difference between 3) and 4), this is explained, *intuitively* at least, by saying that what is being reported in 3) is the illocutionary agent's reliving in memory – his or her memorial re-experiencing as the agent – of the act of closing the door; and in 4) the quite different mental act of perceiving or witnessing this act, as he or she might perceive (i.e. see, hear, etc.) from the outside as it were, a situation in which he or she was not, or had not been involved as the agent (Lyons 1989, p. 176).

Now, the real point is the contrast between (60) and (61). If my intuitions are correct, the contrast is not semantic (as Lyons seems to imply) but pragmatic. It is easy to explain the contrast in terms of pragmatics, as the reflexive is more marked than PRO, and thus tends to trigger M-implicatures, if one listens to Levinson (2000) and Huang (2007). The M-implicature in question is that the perspective from which the action is remembered is complementary to that implied in (60). If the (60) was associated by implicature to an internal dimension, (61) is associated by implicature with an external dimension. In terms of a relevance theory explanation, we can reason on why the more marked lexicalised pronominal is preferred to the null pronominal PRO. Since the more marked item involves greater processing efforts, it needs to be associated with greater contextual effects, such as (I claim) the

complementary interpretation to (61), in particular a not internal dimension (In fact, the external one).

Varela Bravo (1993), in a paper that comments on Lyons' paper, attempts to explain the difference 'I remember closing the door' and 'I remember that I closed the door' through conditions of use, as summed up below:

(I remember) closing the door

1. Acknowledgement of the speaker's communicative intention: I/somebody did something;
2. Evaluation of the action as true/false. That is true: You/somebody did it.
3. Acknowledgment of the action from the point of view of the hearer: Yes, you/somebody did it.
4. Positive/negative evaluation in context: You did well/wrongly.

The utterance would interact with the context and would be functional in the conversational exchange.

(I remember) that I closed the door

1. Acknowledgment of the speaker's communicative intention: something happened:
2. Evaluation of the fact as true/false. That is true.
3. Acknowledgement of the fact from the point of view of the hearer: Yes, that happened.
4. Positive/negative evaluation in context: That was fortunate/unfortunate;

The utterance would interact with the context and would be functional in the conversational exchange.

Varela Bravo basically thinks that an utterance of 'I remembered that I closed the door' presents a fact and focuses on a fact, thus in a tag like 'I remembered that I closed the door. Didn't I?' the pro-verb is 'close' and not 'remember'. Instead, 'I remembered closing the door' focuses on an action done and not on a fact, hence in

the tag question ‘I remembered closing the door. Didn’t I?’ the pro-verb is NOT ‘close’ but ‘remember’.

Varela Bravo makes it appear (albeit he does not use this term) that a (distinct) conventional implicature is triggered by use of each construction. It is easy to see that his considerations give independent support to Harcourt (1999), even if it should be clear that Harcourt and Varela Bravo are making different claims. Harcourt is making a claim about the first-personal mode of presentation, while Varela Bravo is making a claim about whether the context of statements such as ‘John remembers that P’ and ‘John remembers doing X’ restricts the kind of replies one such statement can obtain as a function of the purpose with which the ‘remember’ statement has been made. It is possible that the distinction between a fact being reported and an action being reported is what leads to first-personal interpretations of PRO (in ‘John remembers walking in Oxford’, as, after all, remembering an action requires being involved in the action as an actor (or agent) who has direct access to the action (and its consequences) through consciousness. Of course, the thing remembered in ‘John remembered falling down the stairs’ need not be an action, but merely an event, but even in this case the memory is causally connected to the event and, thus, the experience of the event is somehow involved in the memory. Now this could be a good compromise with Higginbotham’s semantic view, as after all, conventional implicatures are semantic, and not genuinely pragmatic (non-monotonic) phenomena. At the same time, we can see that certain explicatures are nested in the structure available through conventional implicature and these explicatures are a matter of accessing certain more or less stereotypical scenarios.

Before closing this section, I want to consider, albeit briefly ‘imagine’, ‘expect’, ‘dream’ and ‘forget’. For these verbs, I argue that world knowledge is responsible for the explicated content. In fact, the interpretation of the internal dimension of PRO, and in particular the degree of granularity of this internal perspective, depends on the speaker’s and the hearer’s knowledge of the world. In fact, as pointed out by a number of authors, such as Huang (1991), (1994), Clark & Marshall (1981) and Clark & Carlson (1992), Levinson (2000), Blackwell (2000, 2001), Capone (2000, 2001, 2003, 2006), implicatures aimed at enriching utterance interpretations are often determined by the suppositions that are shared by the speaker and the hearer, that is their ‘common ground’.

If I say 'I imagine falling down the stairs' is PRO also associated with an internal dimension? And if I say 'I expect falling down the stairs (if...)', is PRO associated with an internal dimension? What made a semantic association between PRO and an internal dimension in the case of 'remember' was what Shoemaker called the awareness condition. If John remembers falling down the stairs, then he was aware of some experience which caused the memory. 'Remember' also involves a correspondence condition: there must be a correspondence between the event remembered and the event experienced. Now, this compelled me to say that 'Remember' is associated semantically with an internal dimension, but I still proposed a partial pragmatic analysis by saying that the full details of the internal dimension or, to use a terminology from the theory of propositional attitudes, a fully fine-grained internal dimension was provided through explicatures.

Now we have to ask the question whether 'imagine' and 'expect' also involve an awareness condition. If they do not, then the internal dimension of PRO will not be a semantic one, but a pragmatically conveyed aspect of communication. 'Imagining' or 'expecting' seem to be verbs involved in simulating actions or events mentally (to use terminology by Goldman (2006), who explicitly discusses 'imagining' in the context of simulation theories of mind-reading).

When John imagines falling down the stairs, he is probably using information about events happened to someone else in the past. Perhaps he has seen Peter fall down the stairs, and he remembers how Peter felt pain. Thus he may use the information that Peter experienced pain to form a psychological theory about what it feels like to fall down the stairs and concludes that if one falls down the stairs, one experiences pain. So when he imagines falling down the stairs, John runs a simulation of an experience which he saw happen in the past and he recalls that Peter cried, thus evincing pain, and he also has access to a psychology law: if one falls down, one feels pain. As a result of the simulation taking as input a pretend state 'Suppose I fall down the stairs' and some general beliefs, he arrives at the conclusion that he will feel pain. And this conclusion is what authorises us to conclude that when he imagines falling down the stairs, John also imagines feeling pain. The internal dimension is grafted on top of the semantics by pragmatic reasoning. (Of course imagining the pain of an experience has consequences on behaviour, thus RT predicts that the enriched interpretation has greater contextual effects.) Of course John could have run the simulation in a different way. Suppose he is a scientist and

he wants to make some generalizations about the physics of falling down the stairs. Then he is not interested in the internal dimension of the experience. John imagines falling down the stairs having an ulterior purpose in mind, that of simulating a physical experience. Thus the internal dimension is completely missing in this simulation. However, unless aspects of the context do not specify that the simulation is being run for scientific purposes, John will be attributed a state of mind that simulates the internal dimension of one who undergoes that experience, hence pain.

Similar considerations are applicable to ‘John expects falling down the stairs’. The psychological dimension is accentuated, when sorrow rather than pain is involved in the internal dimension as in ‘John expected being sacked’. If what Carruthers (2006) says about mental rehearsal is accepted, John, in expecting to be sacked, rehearses the state ‘being sacked’ and thus produces an emotive response to the situation ‘being sacked’ and this is constitutive of the internal dimension of PRO. But the internal dimension of PRO in the case of ‘expect’ is the result of running a simulation of the simulation John may run of another person’s experience. However, ‘expect’, can also be used in a simulation run for scientific purposes, in which focus is on physics rather than on psychology. Admittedly, such a case is rare, but not impossible. In any case, the internal dimension of expecting something is added only as a result of running a simulation of what it is like to experience the event expected on the part of the person who expects the event, and this is enough to show that the internal dimension of ‘expecting’ is a pragmatic phenomenon.

Two further verbs ought to be considered: ‘dream’ and ‘forget’. ‘Dream’ is in all respects like ‘remember’. If I dreamed murdering Mary, it is implicit that I remember murdering Mary in a dream – hence all considerations I applied to ‘remember’ are applicable to the case at hand. In dreams we usually have sensations in addition to visual images, but it is not clear that the internal dimension has something to do with semantics. In fact, there is no awareness condition or correspondence condition attached to dreaming. It is not the case that if I dream murdering Mary, the dreaming was caused by the awareness of murdering Mary or there was a corresponding state of (my) murdering Mary. All this militates against the internal dimension being incorporated into the semantics. On the contrary, it is reasonable that pragmatics is responsible of the internal dimension. If one dreams murdering Mary, one typically has experiences of fury, sadistic pleasure, contempt for the victim, etc. But this is only part of a typical scenario – it is not impossible to have just visual images with

no accompanying emotions. It is world knowledge that drives the inference, rather than semantics.

Concerning ‘forget’ one can say sentences like ‘I forgot to close the door’, but sentences like ‘I forgot opening the door’ (meaning I forgot the event: opening the door by myself) are weird in English (if we have to follow some comments I found through google). In Italian this type of sentence is fine, but only with a normative interpretation (‘Mi sono dimenticato di chiudere la porta’ (I forgot closing the door) - > I had to close the door but I forgot to do so). The internal dimension of PRO, therefore, even in Italian where it is more certain that this type of sentence is acceptable, is not involved at all.

However, consider cases of ‘forget’ where no PRO is involved. Consider for example the sentence ‘Mary forgot how she felt during her pregnancy’. Here the speaker may very well include both the internal and external dimension among the parts of the event forgotten (There was an internal dimension to the memory of her pregnancy but Mary forgot all details of it). Analogously, if a speaker says ‘John forgot how one feels during an operation, what is at stake is both the internal and the external dimension. However, if one considers the sentence ‘John forgot how he was snoring after the operation’ on the basis of knowledge that John watched the film of the state after the operation, there is no implication of an internal dimension to the memory described as forgotten. All these variations seem to prove that the full internal dimension of memories, knowledge, forgetting, etc. is communicated pragmatically. Consider now constructions with PRO, such as ‘John forgot how to cycle’ or ‘John forgot how to smoke’. There is no internal dimension implication here. (We cannot exclude that one can have corporeal sensations or at least a sensation of happiness when one uses a bike, but it does not appear as though the utterance focuses on these). However, if we change the scenario a bit, as in ‘John forgot how to put up with torture’ the internal dimension is implicated conversationally. I think we can now say that my considerations are not merely tentative, but take into account a range of data. (See Huang 2007 and references therein for the treatment of inferences to stereotype based on scenarios – the operation scenario is Huang’s favourite case).

As a last case, I want to consider ‘know’. The constructions ‘know that’ does not surprisingly does not exhibit the internal-dimension implication (as there is no PRO here), however the constructions involving ‘knowledge-how’ do exhibit the

phenomena noted by Higginbotham (see Stanley & Williamson 2001). Consider initially constructions involving ‘knowledge-how’ without control, with explicit subjects. The sentence ‘John knew how he felt when he was tortured’ certainly implies an internal dimension, however this is far from certain when the sentence is changed a little as in ‘John knew how he was operated’. Suppose that John was totally anaesthetised during the operation. Then it follows that his ‘knowledge how’ cannot take into account an internal dimension – it is a contextual implicature that he knew how the operation was carried out by having watched the video of the operation (I am using a scenario used by Lyons 1989). In the first sentence there is no PRO, yet the internal dimension implication is present. It must be a pragmatic inference in that case. Consider now the sentence ‘John knows how one feels after an operation’. Of course the speaker implicates that John knows both the external and the internal dimension of the feeling. (Yet there is no PRO here). However, if one changes the verb, saying ‘John knows how one sleeps after an operation’ on the basis of the fact that John watched a video of his state of sleeping, there is no internal dimension at stake. (Furthermore, John may know how one sleeps after an operation on the basis of inductive evidence drawn from his having seen many cases of operated patients sleeping after an operation, which would lead him to imagine how he would sleep after an operation (and if he can know how one sleeps after an operation and the kind of problems which operated patients undergo, he would be prepared to pay a private nurse to take care of him, in order to prevent himself from doing harm to his body)). By changing the situation one can cancel the alleged implication

4.6 ‘De se’ and null appositives.

Before ending this chapter I would like to examine, albeit briefly, an example due to Keith Allan (forthcoming)’s important considerations on reference as a pragmatic act. The example was originally ‘The best architect designed this church’, but I want to transform it somewhat, as (62)

(62) Antoni Gaudí believes that he designed this church

Keith Allan argues that strong contextual inferences due to world knowledge will affect the understanding of ‘The best architect designed this church’, and not only

default semantics (Jaszczolt 2005). Could Allan's considerations be transferred to the interpretation of (62)? I would answer this question in the positive, as a pragmatic theory would have to explain not only the fact that (62) has a 'de se' reading, but that on top of the 'de se' reading there is an appositive reading as in (63)

(63) Antoni Gaudí believes that he, who is the best architect, designed this church.

Presumably (63) is obtained for free in an environment in which Gaudí takes for granted that everyone takes for granted that Antoni Gaudí is the best architect. Furthermore, consider the following utterance:

(64)

A: Antoni Gaudí believes that the best architect designed this church

B: Yes, that is true. Gaudí believes that he designed this church.

In this particular context, contextual implications are such that (68b) can be liberated from its first personal interpretation and come to acquire an attributive interpretation. The fact that speaker B agrees with speaker A is strong evidence that speaker B agrees with the speaker's meaning of A's contribution; hence 'he' comes to acquire an attributive interpretation, contrary to the expectation that 'he' is simply a device for reference. The desire to achieve optimal relevance, in terms of Sperber & Wilson (1986, 1995), will compress the attributive interpretation into a pronominal despite the fact that it is semantically implausible that a pronominal should be assigned an attributive interpretation. The way this can be done is that an attributive interpretation can accrue on top of a 'de se' interpretation – in fact, in order to promote optimal relevance the hearer will first of all assign a 'de se' interpretation to (64b) and then on top of that s/he will assign an appositive with attributive meaning as a free enrichment (see Carston (2002), Capone (2008a) and Huang (2007) on free enrichments).

Some readers may complain that I have left undecided the issue whether the inferences associated with 'de se' readings are only implicatures or are, otherwise, explicatures. On an account like Levinson's these pragmatic increments are conversational implicatures. On the account by Sperber & Wilson and Carston, it is tempting to analyse these cases as explicatures. Surely conversational implicatures

and explicatures share the pragmatic derivation. However, I have claimed elsewhere that explicatures are more difficult to cancel – in fact, they may be uncancellable. Surely we can propose that, in so far as they enrich the logical form of a sentence, the inferences associated with ‘de se’ interpretations should count as explicatures. Furthermore, it is clear that the first-personal character of the implicature/explicature is hard to cancel, which supports the idea that we are dealing with explicatures. We are curious to know what happens when we embed the explicature in operators like, e.g. negation., as in

(65) John did not remember walking in Oxford.

It seems to me that (65) preserves the presupposition/implicature that there was an event of walking by himself in Oxford. In other words, the first-personal reading is preserved in the presupposition. In other words, pragmatic intrusion resists despite negation. Analogous considerations apply to (66)

(66) If John remembers walking in Oxford, then I’ll be happy.

As Carston (2002) says, it is a characteristic of intrusive constructions that their explicatures survive embedding in modal operators. If we see that some logical operator has an effect on the inferred interpretation, the inference must have contributed to the truth-conditions of the proposition, as logical operators have scope only on explicitly stated materials (Ariel 2008, 80).

Conclusion

This chapter explores the possibility of deriving ‘de se’ interpretations of pronominals in attitude contexts through pragmatics. After discussing the philosophical literature, by focusing on the tension between a semantic and a pragmatic analysis of de se inferences, I found it fruitful to utilize Harcourt’s idea that ‘de se’ interpretations may involve a mode of presentation like ‘I’ and thus, to respond to potential objections like Feit’s, I revised Higginbotham’s considerations suggesting that the first-personal dimension of PRO in constructions like ‘John remembers walking in Oxford’ should be further characterised by making use of a

mode of presentation like 'I'. Then, in the light of possible objections, I proposed that further shades of meaning could be added by a conversational implicature (explicature). In fact, it is possible that the first-personal interpretation need not use 'I' as a mode of presentation of the reference but some transformation that is suitable in context. Nevertheless, in stereotypical cases, 'I' is implicated as being the mode of presentation used in 'de se' attitudes. The more pragmatic part of the chapter explains why sentences such as 'John believes he went to the cinema' are 'de se' by default, given that 'de se' interpretations entail 'de re' interpretations and, thus, pragmatics promotes the most informative interpretation. The chapter also explains the contrast between 'John remembers going to the cinema' and 'John remembers his going to the cinema' or between 'John remembers going to the cinema' and 'John remembers he himself going to the cinema'. Perhaps the most important part of the paper deals with the internal dimension of 'PRO' and claims that the internal dimension may be the result of a pragmatic, and not of a semantic, inference.

Chapter 11

Immunity to error through misidentification (IEM), ‘de se’ and pragmatic intrusion): a linguistic treatment.

Abstract.

In this chapter, I defend the idea that pragmatic intrusion is involved in ‘de se’ constructions: the ego-concept being a component of the ‘de se’ thought. I defend this idea from a number of objections. I explore the related notion of immunity to error through misidentification and I claim that this too depends on pragmatic intrusion. I defend this view from obvious objections. I take immunity to error through misidentification in ‘de se’ thoughts to depend on the fact that the thinking subject makes an implicit use of the first-person pronominal and there is no question of attributing a referent to the pronominal, since the referent is given in the subject of the thought. In third-person ‘de se’ attributions, some form of simulation can be used to reconstruct the thinking subject as using a form of the first-person pronominal. Immunity to error through misidentification is attributed to the thinking subject through simulation.

Introduction

In this chapter I try to reduce a philosophical topic to a linguistic one. The philosophical topic is ‘Immunity to error through misidentification’ (one who thinks ‘I remember walking in Oxford’ can be mistaken as to whether it was Oxford or Cambridge where he walked, but cannot be mistaken as to whether it was him who walked (in Oxford or Cambridge)) and the linguistic topic is the first-person pronoun (and all its semantico-pragmatic mysteries) in the context of a discussion of pragmatic intrusion (the question whether ‘de se’ thoughts incorporate semantically

or pragmatically the 'Ego' component is an important one). This chapter explores the consequences of accepting certain ideas by Castañeda and after a brief rejection of the attempt by Campbell (1999) to expunge linguistic semantics from the issue of I.E.M. (thus confining it to epistemology), I show in detail that the issue benefits enormously from the findings on pragmatic intrusion. Detailed arguments are provided on why pragmatic intrusion and not merely semantics is involved in understanding the topic of IEM.

PART 1

1.1 Immunity to Error through Misidentification (IEM)

Instead of summarizing the literature on IEM in detail, I will sketch with broad brush strokes the state of the discussion. The positions on IEM are mainly two. There are philosophers who claim that the phenomenon of IEM derives from semantic considerations, mainly considerations of direct reference (Shoemaker (1968), Evans (1982), Strawson (Forthcoming), Peacocke (Forthcoming), Christofidou (2000)) and authors like Recanati (2009) who relate the phenomenon of IEM to the ascription of mental or physical properties either through introspection or proprio-ception or through perception (in the ordinary mode). Shoemaker seems to occupy a unique position in the philosophical panorama since his approach to IEM combines both epistemology and semantic considerations. I will mainly sum up the position taken by Shoemaker and then voice some objections by Wayne Davis (personal communication) which might well be seen as deriving from the alternative camp of philosophers represented eminently by Recanati.

1.2 Shoemaker's position

The article by Shoemaker (1968) is clearly indebted both to Wittgenstein and to Castañeda. From Wittgenstein he takes the notion of IEM and from Castañeda he takes the idea that IEM is linked to use of first-personal pronominals in thought or assertion.

Shoemaker deals with immunity to error through misidentification and seems to imply that this immunity is linked to the first person. Basically, IEM consists in the fact that one cannot ascribe oneself a mental property (e.g. I am in pain) and ask oneself 'Is it me who is in pain?'. In other words, the possibility of misidentifying the subject does not occur. Shoemaker thinks that the first person (the use of the pronominal 'I') has referential uses, a position which, as we will subsequently see, is important for showing where IEM comes from, since in using the pronominal 'I' the subject cannot but refer to himself and also cannot misidentify himself. Shoemaker connects IEM with the pronominal 'I' but he is careful to stress that IEM is not connected with demonstratives in general, as I could say 'That is green' while I am in an hallucinatory state.

What is clear in the discussion by Shoemaker is that the person who attributes himself some property by using the word 'I' cannot fail to identify himself, and this applies both in cases of corrigible or incorrigible statements (Shoemaker distinguishes between incorrigibility and IEM). The subject is available to himself through **introspection** and Shoemaker distinguishes between the ordinary mode of perception and perceptions which we would ordinarily express using the verb 'feeling' by saying that in ordinary perception acts we can perceive the perception act but we cannot do so with perception acts in which feeling is involved. This concept probably needs deepening, but what Shoemaker is after is, presumably, a way to distinguish perception by the senses from introspection, in which the subject is immediately aware of himself (the object is presented to the self from the inside).

I would now like to address a problem noted by Wayne Davis (personal communication) in the formulation of IEM by Shoemaker. According to Shoemaker, at least one aspect of IEM consists in the fact that one cannot ascribe oneself a mental property (e.g. I am in pain) and ask oneself 'Is it me who is in pain?'. The

utterance “I am in pain. Is it me who is in pain?” would be unquestionably odd, but an utterance of “John is in pain. Is it John who is in pain?” is also odd. The oddity of “I am in pain. Is it me who is in pain?” , according to Davis, is not due to the fact that ‘I’ is first-personal and pain mental. It is due to the fact that the first sentence obviously entails the answer to the second. So this cannot be what one means by IEM.

Of course, Davis is right in noting the entailment relationship, which obviously makes such pairs of utterances (in which the preceding utterance entails the following utterance) odd. However, an appropriate question would be to ask how the entailment works when such utterances are embedded in a verb of propositional attitude like ‘believe’. The pair (assertion and question): “John believes he is in pain. Is it John who is in pain?” is not odd, after all. On the contrary, the pair ‘I believe I am in pain. But is it me who is in pain?’ is odd. Suppose Davis is right and the oddity can be imputed to the entailment (or to a pragmatic implication of uttering the preceding utterance, given that now the relationship between the preceding and the following utterance no longer seems to be one of entailment). Then one could replace the pair of utterances related by entailment or a similar relation with a pair of utterances which are no longer related in this way. Suppose I say “I believe I am in pain. But if you were me, could you doubt that it is me who is in pain?”. Now, it is no longer clear that the relationship between the preceding utterance and the answer to the following utterance is one of entailment. The question now seems to be an epistemological/semantic one. One who has such a belief, regardless of the way she herself expresses that belief, cannot reasonably doubt that it is her which is in pain.

Wayne Davis, after rejecting the possibility that IEM consists of the impossibility or implausibility of asking a certain kind of question, after one has expressed a certain belief, says that if IEM consists in the fact that one cannot be mistaken in ascribing to oneself (‘de se’) a mental property. This is the view that beliefs resulting from introspection are infallible. Davis considers cases of people who deceive themselves or people whose desires are unconscious as providing evidence against the view that introspection guarantees infallibility.

However, it should be clear that many of the authors who defend the idea of IEM do not equate it with infallibility; and they do well to do so, because, otherwise, their views could be attacked on grounds similar to Davis’. Immunity to Error through

Misidentification seems to be a logical property of certain kinds of beliefs or statements, rather than only being a claim on epistemology. It seems to me that IEM derives from considerations on what one must accept if one accepts a certain belief which is held in a certain way. I propose that IEM is one of those issues where epistemology depends on the semantico/pragmatic apparatus of language.

Before proceeding with this chapter, I want to deal with a problem in order to restrict its scope. It may appear that I am assuming that a subject's introspective beliefs about his or her own 'de se' mental states cannot be mistaken. This assumption, however, might be questioned on the grounds that it denies that self-deception and unconscious mental states ever occur. Since I do not want to deny that self-deception or unconscious mental states ever occur, it is obvious that my claim concerning IEM must be restricted and qualified in such a way that it does not have this undesired consequence. In other words, we shall retain only the part of the claim that is compatible with the idea that sometimes self-deception occurs. The cases I consider in this paper are clearly not those where self-deception can occur, but those that exclude self-deception. Introspection as a safe method of knowledge of internal states is here seen as confined to those cases where one cannot be mistaken about one's own basic identity ($I = I$). Introspection of the self as self seems to be one case in which self-deception has to be excluded. In this paper I want to tie closely IEM to a semantic claim and I want to make epistemology subordinated to semantics.

Another point I want to make clear from the beginning is that although I assume that 'I' typically refers to the speaker or to the thinker, this need not always be true (a referee mentions arguments by Castañeda), nor is this truth indispensable to my claim about pragmatic intrusion. If my views were expressed through semantic terminology or only through semantic concepts, then this defect could not be easily remedied. But since my chapter is about pragmatic intrusion, all I need for my claims to get through is that the pronominal 'I' be interpreted in those cases when IEM is discussed as first-personal, that is to say as referring to the thinking subject. Surely pragmatic intrusion suffices to guarantee this.

1.3 Semantics or epistemology in IEM? (Campbell 1999).

Since my chapter constitutes a slide from an epistemological to a semantic/pragmatic treatment, before proceeding with my pragmatic treatment I want to present and then refute the position by Campbell (1999), according to which there are no reasons for believing that IEM should rest on a semantic perspective. This section also makes clear how linguistic theses about the first person pronoun could explain why introspective ‘de se’ beliefs cannot be mistaken.

The paper by Campbell links the notion of IEM to the special processes involved in having experiences of mental states, which he takes to be **dedicated** ways of finding out about the world. It appears to me that Campbell is trying to find out whether the notion of IEM can be disentangled from the idea that it depends on the meaning of the first person. He rather seems to opt for the idea that it depends on the nature of the experience of the mental state, on the dedicated process of attributing oneself a mental state. The paper starts with an example, in which IEM seems to be linked with the meaning of an expression. Then he considers a different example invariably showing that IEM can depend on the contingencies of a situation, rather than on the meaning of a linguistic expression.

Campbell considers the case of a businesswoman who has a secretary keeping track of all her movements. The secretary also travels a great deal, but he never has any idea where she is (suppose she herself has a secretary to keep track of her movements for her). We may suppose that the secretary at any moment knows where her boss is, but she knows nothing about where anyone else is, including herself. If you call the secretary, she will always be able to tell you where her boss is, but she cannot give you any information about where anyone else is. You might sometimes get misinformation, but NOT misinformation that will involve a mistake of identification. It cannot happen that the secretary will inform you that her boss is in London, when, in fact, she ought to say that Henry (who is not her boss) is in London (having mistaken Henry for her boss, somehow). As Campbell says, this case of IEM depends on the contingencies of the situation and NOT on the meaning of a linguistic expression.

This “tortured” example was offered as a contrast to the view that IEM can depend on the meaning of the first person pronominal. Unlike other scholars,

Campbell opts for the idea that it is the contingency of the situation of experiencing a mental state that explains IEM. Experiencing or possessing a mental state involves, according to Campbell, a self-attribution. And a self-attribution is a **dedicated** process:

In the case of the first person, what is happening is rather that the subject is using ways of finding out about the world that are, as we might say, “dedicated” to the properties of one particular object, namely that very person. They are not ways of finding out that could equally apply to any of a range of objects. It is for this very reason that although the subject using such a way of finding out can make a mistake, it could not be a mistake about who is in question (Campbell 1999, 95).

Campbell points to the asymmetry between the task of finding out the referent of ‘I’ when ‘I’ is uttered by someone different from the hearer and understanding the use of the first person by the person who uses ‘I’. When you use ‘I’ you do not need any identification act at all.

I should note that I agree with Campbell on the idea that IRM rests on a dedicated process, however I do not agree that this dedicated process is independent of the meaning and interpretation of the pronoun ‘I’. Campbell writes about self-ascription as if it could ever occur without using the word ‘I’, something which I tend to exclude (And Campbell himself avoids the solution that a thinker can make use of the concept ‘whoever is having these experiences’, because this would allow in the case of ‘I am hearing trumpets’ ‘whoever is having this headache’ as subject and one could not ensure identity between “whoever is having this headache and say the person who is experiencing trumpets (see Campbell 1999, p. 94)).

In fact, when Campbell expatiates on the difference between interpreting ‘I’ relative to a context (when ‘I’ is not uttered by the interpreter) and interpreting ‘I’ by the person who uses it, it is clear that matters of interpretation bear on IEM, because Campbell says that the thinker who produces ‘I’ cannot be mistaken about the identity of ‘I’ and I think (and Campbell seems to imply) that IEM is connected with the impossibility of mistaken identity. Of course, one needs to show that IEM derives from semantic considerations. But this is not impossible. Consider the process of

producing 'I' in mental speech? Is it not dedicated too? This process seems to be severely constrained. You do not start with the rule: 'I' refers to the person who produces it and then look for the referent. If you did so, then you could end up retrieving the wrong referent. When you use 'I' you do not go from the context of utterance to see who produced 'I' to the referent of 'I', but you go the other way round, you start with the referent and produce a way of presentation of the referent. So, the process is dedicated in the sense that you cannot be mistaken about the identity of the person producing 'I' simply because the deictic 'I' in silent speech involves a dedicated process: Start with the referent and make sure that the same referent is what your speech is about when you use several occurrences of 'I'. But this looks very much like the description of a semantic phenomenon.

2.1 'de se' in philosophy and linguistics

Immunity to error through misidentification, we have claimed, stems from the use of 'I' and from certain basic semantic characteristics of 'I'. It follows logically that any treatment of 'de se' thoughts incorporating the first person inherits immunity to error through misidentification for the thought, memory or other attitude described.

Whether immunity is direct or derived is not really important for us, since our task is to describe what thoughts a person has in attributing herself a propositional attitude; thus, it goes without saying that if our task is to describe what thoughts are in the mind of a person to whom we ascribe a propositional attitude, then we also have the task of specifying whether those thoughts have IEM or not. In the following section, I discuss both 'de se' thoughts and immunity to error through misidentification. The two issues, in fact, seem considerably entangled. Of particular importance for us is the case of description of a 'de se' thought, in which we attribute a first personal mode of presentation to a thinker/believer and thus we simulate possession of IEM.

In this section I shall present what I take to be the most influential theories on 'de se'. Higginbotham's view is philosophical/linguistic, but I have decided to include it in this philosophical section because it is the only one that has the merit of unifying the first-personal character of 'de se', with phenomena such as the internal dimension of PRO and immunity to error through misidentification. I will mainly use the perspective outlined in Higginbotham (2003), because it is linguistically explicit, in

making recourse to anaphoric concepts and to concepts taken from Fillmore's theory, and I will supplement it with considerations by John Perry (the idea that the essential indexical needs to make use of the concept 'I' at some level of (pragmatic) interpretation).

2.2 Higginbotham (2003).

In his seminal paper, Higginbotham deepens ideas by Castañeda (1966) and Perry (1979). The basic idea is that a 'de se' mode of presentation differs significantly from a 'de re' mode of presentation, having important implications for action and decision. Higginbotham (2003) recognizes that there is something special about first-personal uses of pronominals such as those discussed by Castañeda. The merits of his discussion lie in his pointing out that constructions with PRO may even be more first-personal than uses of 'he himself' and in linking the issue of immunity to error through misidentification to the issue of the internal perspective in connection with PRO (in cases of verbs like 'remember', 'imagine', etc.). He claims that the propositional analysis articulated through the notion of anaphora and thematic roles is superior to the property-based view of Lewis and Chierchia. In fact, according to him, the property-based analysis of beliefs and attitudes 'de se' does not allow the theorist to explain 1) immunity to error through misidentification; 2) the internal dimension of PRO in complements of verbs such as 'remember', or 'imagine'.

Higginbotham accepts Perry's idea that 'de se' attitudes involve a first-personal mode of presentation (involving sometimes the word 'I' or some related notion) and reformulates such a view through considerations based on anaphora and thematic relations.

Higginbotham also accepts Peacocke's (1981) consideration that a 'de se' thought involves the use of a mode of presentation 'self' which only the thinker can use and which nobody else can use in reporting such a thought.

Higginbotham considers cases with PRO such as:

(1) John remembers PRO going to Paris

which is contrasted with (2) and (3)

(2) John remembers that he went to Paris;

(3) John remembers that he himself went to Paris.

The first-personal nature of (1) is expressed through a notation which involves self-reflexive thought:

(4) For $x = \text{John}$, $\exists e$, remember $(x, e, \wedge \exists e': \text{go to Paris } (\sigma(e), e'))$

(1) is different from (2) because it involves an internal dimension. It is the internal dimension which, according to Higginbotham, causes immunity to error through misidentification. We can capture this 'internal dimension' through the expression in logical form of a thematic role: the person who undergoes the action in question. So we can reformulate (1) through (5)

(5) For $X = \text{John}$, $\exists e$, remember $(x, e, \wedge \exists e': \text{go to Paris } (\sigma(e) \ \& \ \theta(e'), e'))$.

With this elucidation in mind, we can explain the following facts:

Only Churchill gave the speech

Churchill remembers giving the speech

:

Only Churchill remembers giving the speech.

Surely someone who listened to the speech remembers that Churchill gave the speech or remembers his giving the speech. But are the speeches which Churchill remembers giving and which another person remembers hearing the same kind of thought? At some level of abstraction they are. At some deeper level, however, they

are not. What validates the inference in the deduction above is the fact that Churchill remembers giving his speech from the inside. So in case he has forgotten giving the speech and someone else informs him that, in fact, he gave the speech, Churchill cannot (truthfully) say that he remembers giving the speech. Memory involves an internal perspective in case PRO is used in the complement clause. Thus, if one remembers falling downstairs, one must certainly have memories of sensations of pain; something which one need not have in case memory is reconstructed through an external narration.

Higginbotham discusses an interesting question. He asks whether mad Heimson who believes that he is Hume has numerically the same belief as Hume. The question, put crudely, is if the belief Heimson has in believing that he himself is Hume is the same which Hume has in believing that he himself is Hume. The answer by Higginbotham is ambivalent. On the one hand, their beliefs are different, so much so that we must say that, in believing he is Hume, Heimson has a false belief while in case Hume believes he is Hume, we shall say that he has a true belief. This is nicely expressed through Higginbotham's anaphoric treatment:

(6) For $x = \text{Heimson}$, $\exists e$, believe $(x, e, \wedge \exists e' (\text{be Hume } (\sigma(e) \& \theta(e'), e'))$.

Since $\sigma(e)$ is anaphorically related to Heimson, there is clearly an **external** component to that thought. However, Higginbotham says that at some level of generality, we can say that Heimson and Hume have the same thought

Be Hume $(\sigma(e) \& \theta(e'), e')$

Higginbotham illustrates this through an analogy to two collapses of bridges. Of course, in one sense two collapses of bridges cannot be the same event, unless the bridges are the same. In another sense, we could say that the collapses of two distinct bridges are the same type of event provided that the bridges have similar characteristics.

PART 2

PRAGMATIC INTRUSION

3.1 Pragmatic treatments

In this section, I will report the pragmatic treatments by Capone (2010b), based on Relevance Theory considerations. Here I spell out certain implications of the ideas in Capone (2010b).

3.2 Capone (2010b) and the pragmatics of ‘de se’.

Capone (2010b) is an eclectic treatment mixing linguistic, cognitive and philosophical considerations in order to predict pragmatic results. His approach is rather eclectic and is a rethinking of pragmatic scales à la Levinson/Horn/Huang in terms of considerations based on Relevance Theory. His ideas, in essentials, are very simple. If one accepts Higginbotham’s considerations on the logical forms of ‘de se’ and ‘de re’ beliefs (to pick up just the most representative of the attitudes), it goes without saying that the logical forms of ‘de se’ beliefs entail the logical forms of ‘de re’ beliefs (See also Tascheck 1985). Hence the possibility of pragmatic scales. On a Relevance Theory line of thinking, the ranking of ‘de se’, ‘de re’ in terms of entailment entails a ranking in terms of informativeness. Then it goes without saying that a ‘de se’ interpretation of a pronominal (in those cases where both interpretations are possible) is informationally richer and, following the Principle of Relevance, greater Cognitive Effects, with a parity of cognitive efforts, are predicted. We may also concoct stories in which a ‘de se’ interpretation leads to some kind of action which the ‘de re’ interpretation would never cause (See Perry; see also Capone 2010b, the pill story). If this line of thought is accepted, then we can explain easily why

(7)John believes he is clever

tends to be associated with a 'de se' interpretation. As Jaszczolt would say, this interpretation tends to be 'default'. Of course, its default status derives from the way the mind is predisposed to calculate inferences and also from the human tendency to standardize or short-circuit familiar inferences that are probabilistically high.

A possible objection to this line of thought may be the following:

You have not given any reason for thinking that the stronger interpretation is as easy to process as the weaker one and no method of measuring cognitive effort is identified. Furthermore, the fact that one interpretation entails another does not ensure that it has greater cognitive effects.

Answer:

While the objection is correct in principle, it does not work in practice for the following reason. A 'de se' interpretation is stronger not only because of entailments, but because of the effects of semantic entailments. A 'de se' interpretation, compared to a 'de re' interpretation' limits considerably the amount of objects that need to be taken into consideration for interpretation. If John believes 'de re' that he is happy, then in principle there is no limit to the objects which we can consider to establish which person John is thinking about (what the belief is about); we surely need a theory of context in order to limit the number of objects that can be taken into consideration. Even if the context is severely restricted, there may at least be a couple of alternatives, such that they would make the interpretation process more costly than in the case of the 'de se' interpretation, which is only about a single object. So it seems to me I have amply shown that processing costs are higher in the case of 'de re' interpretations. I agree that entailments need not be equated with positive cognitive effects, but in the case of a 'de se' interpretation a positive cognitive effect which is a consequence of entailments is that the thought has consequences on the person thought about. So, if I think I have to go, I will go, and I will not refrain from moving paralyzed by the thought that the person I am thinking of could be someone else. So, the fact that a 'de se' interpretation entails a 'de re' interpretation immediately translates into immediate and important positive effects (more or less the kind of effects on action Perry was thinking about in his famous paper).

According to Capone, one may also investigate scales such as the following:

- (8) John wants to go away;
- (9) John wants [him to go away];

- (10) John remembers going away;

- (11) John remembers his going away.

The use of the marked pronominal, instead of less marked PRO, tends to invite an interpretation which is complementary to that associated with PRO. This can be explained in terms of M-scales in the framework of Levinson/Horn/Huang or in terms of cognitive efforts, which tend to pick up an interpretation disjoint from the one associated with the expression involving least amount of cognitive efforts.

Capone also explains certain interesting examples by Perry, which seem to illuminate further the boundary between semantics and pragmatics. Readers are referred for these to Capone (2010b).

Perhaps the most interesting discussion found in Capone (2010b) concerns the internal dimension of PRO, which is connected by Higginbotham to immunity to error through misidentification. Capone argues that, in connection with certain verbs, such as ‘remember’ the internal dimension of PRO is guaranteed by semantic effects up to a certain point, and that at least part of the internal dimension associated with PRO is due to pragmatic effects driven by typical scenarios. With some other verbs, such as ‘expect’, Capone argues that it is less likely that the internal dimension of PRO is a semantic inference and opts for the view that it is a pragmatic increment. Other verbs such as, e.g., ‘knows how’ are examined.

The most radical part of Capone’s ideas is that Higginbotham’s semantic elucidations for verbs such as ‘remember’, ‘imagine’ etc., refined and important

though it is, suffers from a certain weakness, which cannot be remedied semantically, but only pragmatically. Leaving aside formal notation, Higginbotham's treatment of (12)

(12) John remembers walking in Oxford

comes out as 'The agent of the remembering/John remembers that the agent of the walking was walking in Oxford'. But then John should know that he is the agent of the remembering, a grammatical expertise which may not be acquired by anyone at all (See also Davis Forthcoming on this problem). My discussion suggests that Higginbotham's logical form of this sentence could be paraphrased along the lines of: There is a REMEMBERING-event *e* such that (i) the agent of *e* (intuitively, the person who remembers) is John and (ii) the content of *e* (intuitively, what is remembered) is that the agent of *e* participates in an event of walking in Oxford. The main possible argument against Higginbotham's view is that Higginbotham's logical form does not entail that John is aware that he is the agent of the REMEMBERING-event. Higginbotham's logical form thus seems to leave open the possibility, for example, that John, believing that he is Mary, envisions in a kind of third-person way John walking in Oxford, and in actual fact 'John remembers walking in Oxford' seems to exclude this. This argument might have some force (despite the fact that Higginbotham in his paper dismisses this kind of possibility as "pathological") but depends on how we spell out REMEMBER. The possibility I mentioned appears to be pathological because it involves a quasi-memory implanted in somebody else's brain. However, a less pathological possibility is that someone (else) remembers the walking in Oxford (that specific event) in which he was also an agent of the walking (suppose it is also true that he happened to walk in Oxford at the same time when Higginbotham was). Then the only way to prevent this from happening is to set up REMEMBER so that it applies only to events that are remembered from the inside, not from the outside. Now, if this is done through semantics (say a stipulation), full generality is lost, because I claimed in Capone (2010b) that for other verbs also allowing 'de se' interpretations the internal dimension of PRO is not to be specified semantically (due to exceptions). These are the cases of the verbs 'expect' or

‘imagine’. So, instead of opting for a semantic stipulation, I would go for full pragmatic intrusion and also claim that there is a pragmatic presupposition that identifies the agent of the thinking/remembering/etc. event and the agent of the embedded event.

Furthermore, this analysis presupposes that there is a unique thinker of this thought and is incompatible with the possibility that someone else, say God, is having the same thought. (This difficulty was raised by Neil Feit (p.c), and is to be taken seriously). The third kind of problem is that, despite the fact that Higginbotham says that these constructions are first-personal, there is nothing in their logical form that makes them first personal, unless one allows as normative the inference ‘I = the believer of this thought’ (and here endless discussions could arise on how obvious, normative or natural this inference is or should be).

My own view is that the first personal element EGO must somehow be incorporated into the propositional form, NOT at the level of semantics, but at the level of pragmatics. We need pragmatic intrusion – and here the theories due to Levinson (2000), Carston (2002), Sperber and Wilson (1896) come to our aid. Accepting that semantics can be underdetermined, we may incorporate certain elements through pragmatic intrusion. Considerations of parsimony may even lead us to think that pragmatic intrusion, in this case, is to be preferred to incorporation of the component EGO at the level of Higginbotham’s logical form.

Now, I want to dwell on the possible replies to Neil Feit’s objection to Higginbotham. First of all, I sum up Neil Feit’ opinion:

Another reason why I do not think Higginbotham’s account can handle ‘de se’ cases adequately is this. It seems possible that somebody could believe (correctly or mistakenly, it does not matter) that he is not the only thinker of a certain thought, for example he might believe that God is thinking it too. More generally, he might think that he is not the only thinker of any of these thoughts. But, even with this, it seems he could have a ‘de se’ belief. But on Higginbotham’s view - and other similar views – such a belief amounts to ‘the believer of this thought is F’. This cannot be what the belief amounts to however, since he does not think there is a unique believer, the believer of his thought. Moreover, if someone else (God perhaps) really

is having the same thought, then all Higginbotham-style beliefs are false, but he could surely have some true 'de se' beliefs (personal communication in Capone (2010b)).

Now, of course, when I say 'John believes he is not crazy' I do not have in mind other believers of this thought other than John. And, if it is somehow in the background that God and I are the only believers of this thought, I do not think I thereby express or intend to express that John believes that he and God are not crazy. Nor does Higginbotham think so (I assume). The examples by Higginbotham, such as 'John remembers walking in Oxford', are less vulnerable to Feit's objection. Higginbotham's possible reply would be that, given the anaphoric properties of PRO, it goes without saying that the unique believer of this thought (the agent of the remembering) is John and not God (Is not anaphoric coindexation enough to make this clear?). It is not even necessary to resort to the more complicated story that makes the subject of the walking plural: the believers of this thought, assuming a kind of metaphysics in which wherever one is, God is there too. (And I suppose that if talk of God is infused into Higginbotham's story, then certain metaphysical consequences would not be completely absurd).

Of course, the problem raised by Neil Feit becomes more cogent not in the cases of constructions dear to Higginbotham, but in cases of interpretative ambiguities such as (13):

(13) John thinks he is happy.

Here pragmatics is abundantly involved, as even Higginbotham has to admit, and it goes without saying that if Feit's objection has some cogency, this goes up to some point, because if, by pragmatic intrusion, we create an anaphoric identity link between the thinker of this thought and John, the uniqueness condition is valid and thoughts about God's having the same thought are out of the question. (So either we assume that some pragmatic linking between 'John' and 'the believer of this thought' is presupposed, making Feit's considerations otiose, or one needs to insert the anaphoric link explicitly into the semantics).

Another line of response to Feit could be the following. Suppose that the corresponding analysis of John thinks he is not crazy would wrongly imply that John believes that he is the sole individual to have that thought. Even so one could fail to see why this is the case – it would commit one to the existence of a certain thinking event with a unique agent, but certainly there could be other thinking events around with other agents that involve a similar content, and so the agent of that first thinking event could believe that there are.

It does not seem to me that this line of response can dissolve the problem. The problem is that for the sentence ‘John thinks he is happy’, we do not want it to be true if John thinks that God (the thinker of ‘He is happy’) is happy. We want to exclude this case, and to do so we need to add something on top of Higginbotham’s elucidation, presumably the Ego component, which would transform a ‘de re’ thought into a ‘de se’ thought (also along the lines of Reimer 2009, who admits that a ‘de se’ mode of presentation in the presupposition of acquaintance of an object is what distinguishes a ‘de re’ from ‘a ‘de se’ thought).

A ‘precis’ is in order. Despite seeing some problems in Higginbotham’s semantic elucidation, I accept it as the best possible semantic solution to work in tandem with pragmatic intrusion. If pragmatic intrusion is granted, Higginbotham’s semantic elucidation provides very well to provide minimal truth-conditions on top of which full pragmatic intrusion accrues.

4.1 EGO or not EGO?

While Castañeda (1966) in his seminal papers disseminated some really original ideas about ‘de se’ attitudes, and provided the basic examples alighting the theoretical discussion, he was clearly at a fork in having to decide whether ‘he*’ was completely irreducible (a clearly radical and original claim) or whether it could be partially reduced, say by making use of the concept EGO, to appear somehow in the semantic/pragmatic analysis of uses of the essential indexical. The other horn of the dilemma is certainly constituted by Perry’s ideas that beliefs ‘de se’ amount to specifications of mental states in which the concept EGO appears somehow (if it could not be shown to be semantically present, it could be shown to be indispensable for a pragmatic type of analysis). While the considerations by Perry are quite

straightforward and presumably presuppose the at least partial reducibility of ‘de se’ to the EGO concept, Castañeda’s considerations about the irreducibility of ‘de se’ are fully articulated and explicitly deny that recourse to the concept EGO, even if invoked through pragmatic machinery, could be useful. Now, this is unfortunate because there is a clear link between IEM and the EGO component of a ‘de se’ thought.

4.2 Immunity to error through misidentification is the result of pragmatic intrusion

Immunity to error through misidentification is a property which ‘de se’ constructions inherit from the property of the ‘first-person’. However, if my claim that ‘de se’ constructions involve use of an implicit EGO component through pragmatic intrusion is correct, it cannot be true that immunity to error through misidentification is a semantic property of ‘de se’ constructions, although we can legitimately say that it is a pragmatic property of ‘de se’ constructions, being derivative from the EGO component incorporated into ‘de se’ constructions through pragmatic intrusion.

For the time being, we must decide whether we should derive immunity to error through identification from the internal dimension of PRO (or of a ‘de se’ construction) or whether we should derive the internal dimension of PRO (or ‘de se’) from immunity to error through misidentification. This is not a trivial question. We can make this question even more complicated by asking whether the internal dimension is derivable from the implicit use of EGO in ‘de se’ constructions. After all, we could have the following logical chain:

EGO > Internal dimension > immunity to error through misidentification.

If the logical chain above has some validity, and we can establish without doubt that EGO is a pragmatically enriched component of the ‘de se’ construction, then we ‘ipso facto’ show that the internal dimension of ‘de se’ and immunity to error

through misidentification are consequences of pragmatic intrusion and, in particular, the incorporation of EGO in 'de se' constructions.

Have we got independent support for such a line of thought? Recanati (2009) has insisted that not all 'de se' constructions involve immunity to error through misidentification and that in some cases proprioception is involved in guaranteeing immunity to error through misidentification. What is proprioception? While the discussion is undoubtedly complicated, Recanati distinguishes between feeling that something is the case and seeing that something is the case. For example, I can feel that my arm is broken or I can see that my arm is broken. In case I feel that my arm is broken, proprioception is involved and there can be no case for error due to misidentification (it is proprioception that guarantees immunity to error through misidentification). If I see that my arm is broken (say I look at my arm and see an arm that is broken, but I mistake your arm for my arm and I make an identification mistake, then immunity to error through misidentification is not guaranteed). While there is some truth in this discussion, I think it deserves deepening. However, unlike Recanati, instead of placing the burden on the distinction between perception and proprioception, I want to make immunity to error through misidentification depend (at least in basic cases like 'John thinks he is clever') on the awareness of the subject of the thinking experience. Of course, awareness of the subject of experience involves some kind of self-awareness and not proprioception proper or only perception, but the kind of immunity to error through misidentification in cases like 'John thinks he is clever' is different from the cases discussed by Recanati and does not concern objects of experience but subjects of experience. (See also the discussion by Galen Strawson). Thus proprioception may not be the right concept in this case, because it is not the case that the thinking subject is engaged in proprioception in thinking (with some appropriate exceptions, of course: This thought makes me nervous; this thought makes me sad; this thought made me tremble; this thought made me faint). Thinking is the essential relation necessary for establishing a thinking subject. It is the act of thinking that establishes the subject and the identity between the subject of thinking and the subject of the thought. While the person who thinks (14)

(14) I think I am clever

is not particularly engaged in an interpretation process but provides the appropriate EGO concept by the act of thinking and this is enough to ensure immunity to error through misidentification; something different occurs in (15)

(15) John thinks he is clever.

Here the hearer/reader must simulate (as noted by Igor Douven Forthcoming) an act of thinking and in simulating this act of thinking she supplies an EGO concept through inference. Of course, pragmatic inference, utilizing the principle of relevance, independently supports the simulation process and establishes the anaphoric link between John and 'he' and also supplies the EGO concept which is incorporated into the thought by pragmatic enrichment. Having established that John thinks of himself as Ego and that this is guaranteed by the act of thinking in itself, the hearer can simulate John's mental state and, in particular, the internal dimension of the thought (he thinks he is clever or happy because he experiences cleverness or happiness).

In the above, I have proposed that the first-personal dimension of 'de se', as pragmatically implicated, is logically responsible for immunity to error through misidentification (we could also see this case as a case of immunity of error through misidentification being supervenient (in the sense of Chalmers 1996) on the ego-component of 'de se').

If, as I proposed, the EGO component of a 'de se' thought, is due to pragmatic intrusion, immunity to error through misidentification is a consequence of a pragmatic attribution in **reports** of 'de se' thoughts. In naturally occurring 'de se' thoughts which are not reported, it is the act of thinking and the identity between consequential acts of thinking that guarantees the EGO component, and, consequently immunity to error through misidentification.

According to Wayne Davis (p.c.) my thesis is that the semantics together with the pragmatics of first-person pronoun explains why certain beliefs about ourselves are

immune to error through misidentification. This seems incorrect for Davis, as being first-personal is not what makes a belief special; it is being the product of introspection that makes it special. Moreover, it is hard to understand how any fact about the *words we use to express* beliefs can explain why those beliefs cannot be mistaken. How can facts about the English word 'I' – a mere letter or speech sound with a meaning and use – explain why a certain propositional attitude has a special epistemological status? Propositional sentences are not sentences or strings of letters or speech sounds. (Notice, however, that I am not claiming that IEM is ONLY a semantic phenomenon but that it is an epistemological/semantic/pragmatic phenomenon).

Ok. Suppose that introspection is, as Wayne Davis proposes, the essence of IEM. Is there a way to argue that introspection is not really possible (everyone should accept at least that it is facilitated by the use of the first-personal pronominal) unless a predicate is assigned to a subject and that introspection specifically necessitates assigning a predicate to a first-personal subject? The role played by semantics is evident. But this does not amount to saying that the first-personal pronominal is essential in self-attribution of an introspective thought; surely the fact that the thinker and the attributor are the same makes it possible that no mistake is made in the attribution even if no first-personal pronominal is used; however, in expressing that thought to others it is necessary to use the first-personal pronominal.

A related question is whether all beliefs and thoughts have to be expressed in words. As Davis (p.c.) says, a belief or thought about ourselves is distinct from any words we use to express it. The belief that I myself am in pain is very different from the sentence 'I am in pain'; the belief does not contain four words, is not a sentence of English, cannot be seen, heard, written or uttered, and so on. The sentence may occur in inner speech, but inner speech is not belief or thought. An I-thought does not contain the word 'I'.

Ok. Let us grant that at least sometimes one need not use language to form a thought, in particular a thought about oneself. Let us suppose that this thought is not very different from a sensation (pain). The fact that pain occurs in myself is sufficient to make the case that a mistaken identification cannot occur. Suppose then that it is the object (my body) which prevents me from making a mistaken identification. Yet I suppose that the body is at least part of myself and, thus, the

thought, if it has to be expressed, requires the logical subject ‘I’ and only this, because a definite expression could not preclude a mistaken identification. So, the argument is not that a first-personal thought requires the word ‘I’, but that it requires it if it has to be expressed. The question of misidentification, in a thought which is sufficiently similar to the attribution of a sensation, cannot occur; but it can occur when the thought is expressed. This is all we need. So I may concede that Wayne Davis is partially right about belief – even if I should say that the kind of belief he has in mind is more similar to sensation than to the propositional attitudes we normally discuss.

4.3 Why immunity to error through misidentification is logically INDEPENDENT of the internal dimension of PRO/’de se’.

Admittedly, the reasons which I furnish in this section against making a logical connection between the internal dimension of PRO/’de se’ and immunity to error through misidentification depend on some previous considerations of the inferential behavior of de se/PRO, presented in Capone (2010b). There I wanted to make the provision/expression of the internal dimension of PRO/’de se’ a pragmatic constituent of the thought. However, after some discussion I moved towards the more balanced view that, in general, especially with verbs such as ‘remember’, the internal dimension of PRO is semantically associated with the specific construction (PRO, in our case) (See Ninan 2009 for an important paper on imagining from the inside, which amounts to imagining a scenario from the point of view of one of its participants, in a centred world). Then I have speculated that the internal dimension (constituent) supplied through the semantics is only partial or gappy (in line with views by Carston (2002) on semantic underdetermination) and that pragmatics is responsible in part for supplying a partial internal dimension. For certain other verbs, such as ‘expect’, ‘know how’, etc. I have speculated that the internal dimension constituent is fully provided through pragmatics.

Now, what are the consequences of the acceptance of the views above for the plausibility of the view that immunity to error through misidentification depends on the internal dimension of PRO/de se? The most immediate consequence would be that, in the most straightforward cases, like ‘**expect**’, or ‘**imagine**’ ‘de se’

constructions (he* or she*) should not be associated with immunity to error through misidentification, if IEM depends on the properties of the internal dimension through semantics. Thus, someone who expects to leave for Rome tomorrow may legitimately have some doubts as to whether he himself is involved in the thought that he will leave for Rome tomorrow. But this is absurd. Immunity to error through misidentification must be granted for cases such as ‘expect’ and ‘imagine’ as well and this shows that immunity is not logically dependent on the internal dimension (which is implicated in these cases, if my view in Capone (2010b) is correct).

What other consequences follow from the fact that the internal dimension of PRO/‘de se’ is only partially semantically expressed and partially pragmatically articulated in cases such as ‘remember’? If we grant the logical dependence between immunity to error through misidentification and the internal dimension of PRO, we paradoxically arrive at the result that the greater the pragmatic enrichment in connection with the internal dimension of PRO/‘de se’, the greater the immunity to error through misidentification. However, nobody says or is willing to accept that immunity to error through misidentification is a gradable notion.

My initial idea that the internal dimension of PRO/‘de se’ may serve to support immunity to error through misidentification and not to establish it is not implausible; however, the immediate consequence of having such a view is that the internal dimension of PRO is useful in establishing immunity to error through misidentification only in those cases where there can be some doubt because a sentence is ambiguous. Consider, again an ambiguous sentence similar to one example by Recanati:

(16) He thought his legs were crossed.

Depending on whether he was only seeing his legs crossed or was also feeling them (proprioception being involved), (37) presents (or does not) a case of immunity to error through misidentification. The internal dimension of the pronominal ‘his’ is clearly projected through a pragmatic enrichment and, thus, proprioception is responsible for promoting immunity to error through misidentification through the

intermediation of a first-personal pronoun. The pragmatically enriched internal dimension and proprioception go hand in hand and serve to reinforce immunity to error to misidentification in the sense of disambiguating a sentence which is interpretatively ambiguous.

4.4 Wayne Davis (Forthcoming), 'de se' and indexicals.

In this short section, I make connections between IEM and Davis' view that 'de se' is linked with deixis. Davis thinks (in essentials) that 'de se' attitudes are to be explained by reference to deictic concepts. The thinking subject thinks of himself through a deictic. This is similar to what I have claimed myself, although Davis is more detailed. I was content with an 'I' concept, while Davis distinguishes between a deictic, a demonstrative and an anaphoric use of 'I'. The deictic use of 'I' is probably what is involved in 'de se' thoughts, deictic uses being licensed by what Davis calls 'presentations'. The thinker thinks of himself and has a presentation of himself that gives interpretation to his use (whether mental or verbal) of 'I'.

I depart from Davis in recognizing a dichotomy between the **thinker's** use of 'I' in thought, and the **hearer's** interpretation of 'I' or 'he*' in an ascription of thought. The thinker's use of 'I' in thought needs no special act of interpretation and involves immunity to error through misidentification in that no identity is needed or established, as there is no interpretation problem from the point of view of the thinker, who surely has a presentation of himself which is perhaps tacit and who keeps track of himself and his identity through the act of thinking, rather than through the act of interpretation. Instead, the ascription of 'de se' attitudes (to someone else) involves an interpretation problem and keeping track of the referent and mode of presentation used by the thinker either through a simulation or through a pragmatic act of interpretation guided through the Principle of Relevance or both. The two perspectives are different and surely the use of 'I' in ascription of 'de se' attitudes involves both an internal and an external dimension. The deictic use discussed by Davis may be suitable to both dimensions, provided that we are clear that a 'presentation' or 'self-presentation' is involved in the thinker's awareness of EGO, while a simulation or pragmatic interpretation is involved in understanding the presentation which the thinking subject experiences. Perhaps it would not even be

incorrect to say that we can speak of a deictic use when referring to the hearer's interpretation problem, while from the point of view of the thinker there is no interpretation problem and thus it is not a matter of establishing the content of the deictic thanks to contextual coordinates. All that is required is the thinking act and the thinking act is its own context and also its own content.

Before closing this section, should we be content with Davis's exposition? While surely Davis' story resolves the problems he himself raised in connection with Higginbotham's theory (along the lines of the problems I myself discussed), he does so in an ambiguous way. Is the use of the deictic a semantic or a pragmatic component? I was clear that pragmatics was involved in establishing the EGO concept in 'de se' attitudes – even in cases of PRO, which are particularly problematic for Davis since PRO does not receive content from a context and thus is not easily assimilable to a first-personal concept). If we accept the considerations by Davis, we should have a double interpretation process. The provision of an EGO concept and, then, the interpretation relevant to a context of use (but this I admitted through lavish use of anaphora). From the point of view of the thinking subject, however, there is no pragmatics, since he has direct introspective access to his/her own thoughts. Pragmatics is involved only from a third person perspective, that of a hearer who tries to reconstruct the speaker's thoughts and self-awareness.

5.1. 'De se' and modularity of mind: cancellability?

Finally it is time to examine the issue of the cancellability or non-cancellability of the 'de se' inferences I have discussed at length. Non-cancellability *per se*, as Grice was well aware, does not militate against the pragmatic nature of an inference. I have claimed elsewhere that explicatures are non-cancellable and the motivation I gave for this is that explicatures tend to be motivated by problems in the logical form, when a sentence is perceived to be blatantly false or a logical absurdity and pragmatics is there to help and remedy the problem. Since the explicature is the *Deus ex machina* of the semantics, I have claimed in a number of publications that it is and should be non-cancellable. This seems to fit in with a modular view of pragmatic processes, as argued in a number of publications. (See Capone 2010b, Capone 2011 for detailed arguments).

We saw that the incorporation of the Ego concept was the Deus-ex-Machina of the semantic treatment à la Higginbotham, protecting this treatment from all the objections raised by e.g. Davis (Fortcoming).

One further reason for opting for a pragmatic treatment of ‘de se’ is, of course, the parsimony of levels that it affords us, as one can eliminate at least an important meaning component from the semantics, obtaining it for free from pragmatics.

One last reason for opting for a pragmatic level of meaning in ‘de se’ attributions is the differential mechanisms of ‘de se’ thoughts in view of what happens in the mental processes of the thinker and of what happens in the mental processes of the hearer. The hearer is in a different position, both with respect to calculation of the Ego component and of the anaphoric links within the ‘de se’ ascription and with respect to the attribution of Immunity to Error through Misidentification. The disparity between the position of the thinker and the position of the speaker/hearer in connection with pragmatic inferences was noted in an article by Jeff Speaks (2006). The disparity between the thinker and the speaker/hearers’ stance to the inference is due to the fact that **luminosity** is available in thought, introspection being a guide to one’s intended meanings, while the meanings projected by the speaker and understood by the hearer in conversation do not rely on luminosity but on an explicit effort to get intentions across through contextual clues and cues.

While immunity to error through misidentification is presupposed for the thinker in virtue of the continuity afforded by the act of thinking (thus immunity seems to be an ‘**a priori**’ category of first-personal thought, something similar to a pragmatic presupposition along the lines of Maier 2009) and by the fact that in thinking the question of misidentification cannot arise; for the hearer, immunity is a logical consequence of the pragmatic inference involved in assigning an ego component to the ‘de se’ thought. Simulation and, also pragmatic interpretation flowing from the Principle of relevance are clearly involved.

The disparity between the speaker’s perception of himself as himself and the hearer’s ascription of ‘ego’ to the thinker has interesting consequences concerning cancellability. The speaker’s perception of himself as himself is clearly non-cancellable; the hearer’s ascription of EGO to the thinker of the ‘de se’ thought is driven by contextual clues which lead the interpretation process in a certain direction,

from which it is impossible to go back. So both from the speaker's and the hearer's perspectives it is impossible to cancel the EGO component of the thought.

Implicitly, I have replied to some qualms by Coliva (2003) about the idea that Immunity to Error through Misidentification depends on the ego concept incorporated in 'de se' attitudes. Her main objection to this idea is that the use of 'I' in 'de se' thoughts (whether explicit or implicit) is not enough to guarantee a first-personal thought. Coliva speaks of the split between speaker's reference and semantic reference (see also Jaszczolt Forthcoming for a similar argument). Given this split (which has emerged especially in the discussions of Donnellan's attributive/referential distinction), it may not be correct to say that immunity to error through misidentification depends on the presence of a pronominal like 'I' in logical form. The case discussed by Bezuidenhout (1997) (Bill Clinton: The Founding Fathers invested me with the power to appoint Supreme Court justices) does justice to the ideas and doubts exposed by Coliva. In the example by Bezuidenhout 'me' is used attributively, and not referentially. Of course Coliva does well to address the issue of the pragmatic nature of the incorporation of the EGO-component in 'de se' attitudes. However, we get the impression that her skepticism on the idea of deriving immunity to error through misidentification is not completely justified, given the heavy presence of pragmatic intrusion in propositional forms. Given the non-cancellable character of the pragmatic inference which I posited in 'de se' thoughts, the fact that 'I' can be interpreted attributively, rather than referentially is not a problem. Of course, my claim that immunity to error through misidentification follows from the Ego-like nature of 'de se' should be confined to cases where Ego is interpreted referentially. But this is, of course, presupposed by the 'de se' semantic/pragmatic analysis. Again, we should distinguish between the interpretation of the construction (e.g. I believe I am happy) on the part of the speaker, which heavily relies on Mentalese (the speaker has direct access to his own thoughts, and, thus, ego as used in 'de se' constructions is clearly and directly referential) and the interpretation on the part of the hearer. When we examine the dimension of the hearer, we see that the interpretative problem of 'de se' constructions consists in assigning, through pragmatics, an inferential increment making the logical form more plausible than it would otherwise be. The pragmatic enrichment, thus, could not make use of an uninterpreted EGO component, but has to make use of an interpreted EGO component, a component that is referential and not attributive (as in the

example by Bezuidenhout). Of course, if we accepted a view in which the EGO component were assigned at the level of the semantics (say by identifying PRO with 'I' or an EGO-concept), then Coliva's objections could be certainly and dramatically applicable. But this is one more reason for opting for a semantico/pragmatic treatment, rather than for opting for a semantic treatment only. In a sense, we owe to Coliva the intuition that pragmatic intrusion resolves problems that would otherwise be insuperable.

Before bringing this chapter to conclusion, I want to point out that my considerations have something in common with Maier (2009). Maier essentially distinguishes between 'de re' and 'de dicto' attitudes deriving 'de re' cases by accommodating presuppositions at top level in the DR. 'De se' cases are clearly cases in which the object one is acquainted with is a centred object. Only the presuppositions of discourse can establish whether the object one is acquainted with is a centred object. Now, in my discussion I have made an effort to disentangle cases where the presuppositions of discourse are responsible for IEM and the cases where the hearer needs to undertake an inferential process to simulate the mental processes of the person whose thoughts are being described in an assertion.

Conclusion

This chapter has been loaded with theoretical considerations and their consequences. Presumably we have reached a stage in which, pragmatics, which originated in philosophy and was propagated outside philosophy giving impetus to communication-oriented linguistic views, can serve to throw light on philosophical topics. I cannot exaggerate the importance of seeing the phenomenon of immunity to error through misidentification as a consequence of pragmatic intrusion. It is true; I have reached a stage in which the theory has become loaded with various consequences of previously accepted views. However, it is the nature of interconnected considerations and interlocking ideas one finds in this chapter, that makes it rich, by provoking novel and perhaps radical discussions of a phenomenon of which we knew little, before putting some thought to pragmatic intrusion.

Chapter 12

The pragmatics of quotation, explicatures and modularity of mind.

Abstract

This chapter presents a purely pragmatic account of quotation and argues that it is able to accommodate all relevant linguistic phenomena. Given that it is more parsimonious to explain the data only by reference to pragmatic principles than to explain them by reference to both pragmatic and semantic principles, as is common in the literature, I conclude that the account of quotation I present is to be preferred to the more standard accounts. Alternative theories of quotation are treated in an intellectually honest way.

Introduction

This chapter presents a purely pragmatic account of quotation and argues that it is able to accommodate all relevant linguistic phenomena¹. Given that it is more

¹ I assume that this chapter, which is an exercise in radical pragmatics, draws much from Recanati's ideas and pushes them to extreme consequences (I have no idea whether Recanati would approve of what I have done; but I am committed to the idea that this is an off-shoot of his theory of quotation). I also assume that this paper is in line with the idea by Recanati (2004) according to which the main role of pragmatics is to boost semantics by amplifying the expressive potential of human languages. This amplification power results from the fact that pragmatics can use already existing linguistic resources to create new ones on the basis of a limited set of pragmatic principles. Like Saka (2005) I will focus on the role played by the speaker's intentions in quotation practices and I will reject the idea that a pragmatic treatment of quotation is to be equated with Humpty-Dumptyism, given that intentions depend on expectations based on internalized norms and on general language-transcendent principles of cognition (Saka 2005, 5). My view of quotation will be compared with my views of pragmatic intrusion and explicatures (to anticipate, I consider the effects of quotation as products of pragmatic intrusion into truth-conditional content). I will argue that the features of my theory of quotation and the features of my theory of explicatures fit like plugs in sockets and such a fit serves to render both theories stronger. I take the interlocking nature of two modules of the same theory (or of a wider theory) to provide greater confirmation for the

parsimonious to explain the data only by reference to pragmatic principles than to explain them by reference to both pragmatic and semantic principles, as is common in the literature, I conclude that the account of quotation I present is to be preferred to the more standard accounts. Alternative theories of quotation are discussed in a critical way.

This chapter has the following structure. In section 1, I address in a general way the problems which a pragmatic theory of quotation must tackle: the asymmetry between written and oral language; the possibility of quoting a sentence originally uttered in a different language; the fact that in the oral language there is nothing that can be demonstrated through a demonstrative pronominal; the fact that we can master the practice of quotation in languages different from our own; and the fact that we can learn the quotation practice without formal instruction.

In section 2, I discuss the standard theories of quotation (including some old ones like the demonstrative theory) and after discussion of the disadvantages of these theories, I focus on Recanati's theory, which opens the way to a radical pragmatic perspective. I also consider mixed quotation and propose that pragmatic intrusion can happily deal with it. In section 3.1, I connect the discussion of quotation with the discussion of cancellability of explicatures. In section 3.2, I discuss certain differences between Saka's view and my own view, which depend on my focus on interpretation problems. In section 3.3, I deal with camouflaged quotations and I argue that they can be dealt with in terms of pragmatic intrusion. In section 3.4, I address a crucial problem posed by Predelli (2003) to a radical pragmatic theory of quotation.

validity of the two modules in question. If we can show that there are other modules (of the same theory) that have interlocking capacities, we will be able to provide greater confirmation for each module individually. So far my considerations on explicatures tie in with my considerations on Grice's circle, with those of modularity of mind, with those on the attributive/referential distinction and with those on quotation.

1. On Quotation.

Quotation is an important practice, because it allows one to separate someone's voice from the reporter's voice in a direct or indirect report (where mixed quotation occurs). In fact, in direct quotation, we attribute words as well as thoughts to other people. Given that in quotation what is most important is to separate one's voice from another, we predict that that quotation is not merely a semantic phenomenon, given the availability of pragmatic resources that suggest the speaker's differentiation between one's voice and another. Consider a sentence like (1), where English words are ascribed to a non-English speaker:

(1) Galileo said, "The earth moves!"

Such uses, accepted by Sperber & Wilson (1986, 228) and by Recanati (2010, 223), argue in favor of a radical pragmatic view of quotation. In fact, given that the hearer can infer that the quoted item was uttered in a language different from English, it is evident that pragmatic resources must be involved in quotation. From the very beginning, I should point out that although quotation is standardly taken to involve quotation marks, this must be regarded as the result of a focus on the written language; in fact, there can be more than one way of signaling quotation, and in the spoken language one can use certain devices that distinguish between voices for example by showing the different qualities of the voices. This is one of the many resources available which can indicate that a speaker is quoting an expression or an utterance or a portion of an utterance.

This practice of quotation, in written English, consists in putting a phrase or a sentence within quotation marks. However, it is accepted practice not to use any symbols for quotations in logic textbooks (for those cases in which it is clear that one is discussing the forms of words, and not their extension). What happens in spoken language is even more liberal (Washington 1992). It may happen that a philosopher who reads aloud a passage from his article will be at pains to say 'Quote', then the quoted words, then 'Unquote'; and it may happen that sometimes we signal through finger dancing that quotation marks are to be perceived somehow. However, in the remaining cases, most people fail to use quotation marks in the spoken language, since it is clear in context what the intended meaning is, whether one is talking about words or about a word's extension.

Probably the most thorny cases for philosophical theories of quotation are quotational devices such as italics. If one assumes that a demonstrative element is present in the sentence, then surely the italicized part cannot count as the demonstrative element (the demonstrative element is perfectly explicit in Davidson's analysis) in, e.g. Davidson's analysis of quotation. (I suppose that one reason for objecting to the view that italics can be used as a demonstrative element is that a demonstrative is normally an entity which can be distinguished from the thing demonstrated; in the case of italics it is not easy to distinguish a demonstrative element, as the words mentioned and the demonstrative element should be fused together, which is like using a road sign that points to itself say in virtue of changing the color of the sign. Why is it that a road sign like this does not exist? It does not exist because we take the road sign to point to something external to it and not to something inside itself². This is presumably what is behind the idea of demonstrating something. Surely we could use a road sign in this way, by relying on some marked characteristics to point to the fact that the authority who made use of the road sign wanted it to point to itself; but this would make the use of the road sign difficult to understand, because the point of a road sign is to demonstrate something outside itself; once we become used to the demonstrative practice, it would be difficult to modify it by adding peculiarities to it. To return to italics, italics can show that an expression is being mentioned but it is not a way of pointing to the expression). At this point, it may be convenient to distinguish between using a demonstrative and demonstrating (as Reimer 1996 does). While it could be argued, somehow, that even

² Paolo Leonardi (personal communication) has got a different but similar example: you arrive close to Rome and at the entrance of the city there is a road sign with the written string: ROME. Here the road sign indicates the border between Rome and the rest of the territory. A reviewer also point out that we must compare the P with an arrow on streets in cities. Horizontal arrow tells us to take P as quoting ('There is a parking there'). A vertical arrow tells us that P is an injunction /prohibition /limitation, not a quote. In Finnish we use case forms like the following: Nominative metsä vs. Partitive metsää ('forest') e.g. 'Valtion metsä', with straight arrow, means: "There is a State Forest there" (in direction of arrow). 'Valtion metsää' with no arrow means "This is all a State Forest". This is used to express a definition, no quotation is involved here.

the use of italics is a way of demonstrating an object (say, a word), surely it cannot be granted that italics amounts to usage of a demonstrative pronominal, as such a pronominal would have to be external to the object demonstrated. I take it to be a universal of language use that a demonstrative can be used to point to anything but to itself; of course there may be apparent difficulties with this, such as cases like “This sentence is blue”. Here the demonstrative refers to itself (as well as to three other words). Things are not so easy though. I can say “This sentence is blue” by inviting the hearer to look at the way the sentence is written (a written token of the sentence, actually) and I am inviting the hearer to consider the whole sentence and not just the demonstrative; it happens that the demonstrative is part of the sentence, but it cannot be used to demonstrate itself in isolation from the sentence in which it occurs; this is a clear limitation to the idea that a demonstrative can demonstrate itself. Paolo Leonardi (personal communication) considers that there are more complicated cases like ‘This utterance is false’ in which ‘this’ stands for the entire utterance; or there is the case in which you ask me to point to an Italian demonstrative and I answer ‘This’. As to the second case, I would be happy to consider it a compressed use; what the speaker means is: This is it: ‘this’. As to the former case, I would like to notice that at least with some demonstratives the self-referential use is not possible, as in ‘That utterance is false’, which means that it may be legitimate to claim that ‘This utterance is false’ is not a real demonstrative but is a use that is parasitic on demonstratives. ‘This utterance is false’ means ‘The utterance I have just proffered is false’; but now it is clear that ‘This’ in this case is nothing more than a definite article in combination with an implicit description.

Now, even if we concede that italics could be used (pragmatically?) to demonstrate an object, its demonstrative power can never amount to the full power of a full pronominal, given that it would necessarily have to point to a demonstrated object but it could not point to an italicized object. However, Saka (personal communication) says that this seems to be false. Consider: “*Blue* is in italics”. But Saka’s example makes explicit use of the word ‘italics’ and it is this word which helps italics to demonstrate the expression in italics, italics being part of the mentioned token. Without the explicit use of the words ‘in italics’ the hearer would be in the dark as to whether the speaker intends italics to be part of the quotation or not. It appears that, when no explicit words accompany such uses, it is the speaker’s intentions which make matters clear.

Suppose one accepts that quotation marks are a grammatical device, one that is used in virtue of knowing a convention about language (and about a particular language). How does it come about that we are mostly able to understand how quotation works in languages which are not related to ours? We can interpret quotation even if we are not fully aware or have never been formally instructed that a certain grammatical operation is needed to obtain quotation. This seems to point to pragmatic theories of quotation.

Wayne Davis (personal communication) insists that some linguistic conventions are common to many languages, especially written language conventions. One group of examples are the conventions governing punctuation marks like ‘?’ and the period; quotation marks falling in this very group. Another are Arabic numerals and mathematical function symbols like ‘+’. Second, the interpretation of quotation marks is such that they can be used to refer to the words of any language; there is no reason why speakers of all written languages cannot use quotation marks with the same interpretation. I do not quarrel with this idea that there are some constant ways of expressing quotations in the world languages. My point is merely that if quotation marks were expressed through a minimally different usage in a different language, second language users could be capable of grasping the quotational meaning without formal instruction³. If conventions are, usually, of an arbitrary kind, a point which even Davis must grant, then we would expect that a second language user, faced with an expression such as ‘das auto’ (‘das auto’ being a foreign word for him), would pause to ask himself if ‘.....’ is a quotation in the new language, for, for all he knows, given that every language has its own conventions, ‘...’ could be used for a different purpose. But, in fact, second language users do not hesitate like this in using quotation, which strongly points to a pragmatic, rather than to a semantic, conventional, mechanism.

Another quibble may derive from my use of ‘formal instruction’. Quotation may be a conventional mechanism, even if it is not acquired by formal instruction. At this point, we may hesitate. Was it a formal instruction or not, when our teacher corrected

³ I was told in conversation that German and French have different conventions about quotation. Of course, I do not exclude that conversational inferences can get grammaticalized in a language and I admit it is possible that they get grammaticalized in a different way. However, the pragmatic principles underlying such inferences may be the same.

our language use, when we forgot to use quotation devices? The stress by Chomskyan linguistics on language emerging naturally has made us forget how much work our (old) teachers made in correcting our copybooks and marking them in red. I assume that was part of a formal instruction, attesting to the fact that many things about language use were learned in a social context. But now my opponent may want to argue that, in fact, quotation devices are conventions learned through formal instruction. My answer is: not necessarily. The teacher may well reinforce some principles of language use. In fact, he may never say that this is a convention of Italian or of English. It is possible, in theory, that the quotation mechanism was learned as a convention, one of those tacit conventions which one never bothers to explain through formal instruction, belonging to those things which are too obvious to be explained or learned through a formal instruction. In this case, it could be like one of the many grammar rules which we learn without anyone having to point to them. But then in this case, we would have to have a rule for the oral language and a rule for the written language. And for foreign languages? The fact that we have internalized a rule or a set of rules for our own language makes us ready for the expectation that a different language may have different mechanisms for quotation. And yet, when encountering a foreign utterance, we have no such expectation, but on the contrary we go on behaving as if a universal language were available. Is not this a reason to suspect that such a universal of language use must come from pragmatics? (Alternatively, we must explore the possibility of a grammatical principle like binding Principle A or B or C).

Suppose there was no convention in the language for signaling quotation. Suppose nothing was really available to point to quotational uses. Nevertheless, it is a universal of human languages that a language must be able to express somehow the cognitive operation ‘quotation’ and if there are no explicit resources, then it must use resources which are not conventional⁴. Presumably, there is a tendency in human languages to develop grammatical resources that make the quotation operation

⁴ I received the following comment: “Perhaps it is, as a matter of fact, universally true that languages are able to express that operation. It is not so clear to me why they *must* have it. Is there a principled reason to think that paraphrase could not serve the same cognitive purposes that quotation is thought to serve?”. This involves acceptance that quotation and paraphrase might be functionally equivalent. However, without quotation it is hard to explain what a lexical item means; it is at least reasonable to accept that every language must have devices for teaching and learning lexical items, quotation being one of them.

available on the basis of pragmatic resources. The clear-cut asymmetry between formal quotation operations in the written language and pragmatic quotation operations in oral languages attests that pragmatics is somehow responsible for quotation and that pragmatic intrusion is necessary. (For additional arguments that pragmatic quotation precedes formal quotation, see Saka 2006, section IV.2.)

Wayne Davis (personal communication) opposes this view:

“If by ‘explicit resources’ you mean morphological or punctuation conventions, then I think your claim is clearly false. In particular, I believe there are conventions in relation to intonation and phrasing that function like quotation in speech. Contrast the way we say *John said a word* (meaning that John uttered at least one word) with the way we say *John said ‘a word’*.”

Now, I quite agree with Davis that, in the oral language, we find resources allowing us to point to a self-referential function. However, one must explain cases like *John is ‘very happy’*, where there can be no colons parasitic on the written language and where one must find some implicit resources to signal that the speaker is mimicking John’s utterance (or voice). Would Davis say that conventions of language use are available in such cases? It is a pity, for such a view, that one has many ways of marking John’s utterance as belonging to John: one can use a voice that mimics John’s voice, one can use facial expression to dissociate oneself from that portion of the utterance, one can use different quality of voice to separate various portions of the utterance, etc. Is there a limit to what one can do to contrast the two portions of the utterance? Provided that the contrast is made clear, one can choose whatever is more congenial to oneself.

It can be argued that it is convenient to say that the cognitive operation ‘quotation’ is, after all, always pragmatic, and that written language has grammaticalized resources otherwise available through pragmatics. It is not impossible to argue that quotation, even in the case of written languages, is a pragmatically effected operation. If there is such a possibility, then the asymmetry between written and oral uses is eliminated in one stroke. Everything else being the same, one would tend to choose a theory that is more parsimonious and tends to reduce levels of meaning (following Modified Occam’s Razor). Of course, it is possible that the simplest theory will not suffice to explain most details and data; in

which case, this possibility ought to be reconsidered. I hope this will not happen in this case.

Before launching such a reductive and ‘radical’ enterprise, I want to ponder in a preliminary way on whether a theory of explicature can assist our understanding of quotation. If we take Recanati’s words, a theory of explicature can explain many of the subtleties of quotation. What we should be clear about, however, is not whether some feature of the theory of explicature can help our understanding of quotation, but whether the theory of explicatures can be applied completely to quotation. Explicatures deal with intuitions about truth. Is truth at stake when quotation is concerned? Surely, if one utters (2)

(2) ‘Time’ has four letters

The utterance can be said to be true if one attributes the predication ‘has four letters’ to a particular expression type. It can be said to be false if, despite quotation marks, it is interpreted as attributing the predication to the metaphysical entity Time.

So, in a preliminary way, we have established that questions of truth and falsity are at stake in the case of quotations and, thus, explicature is a potentially relevant notion that offers us a promise to tame the mysterious offshoots of quotation theory. I should also say, in a preliminary way, that I will try to resolve some puzzles which beset Recanati’s application of the notion of explicature to the issue of quotation.

2. Overview of theories of quotation.

In this section, I shall succinctly discuss the major theories on quotation. After discussing the defects of the old theories, I focus on Recanati’s view, which opens the way to a radical pragmatic treatment of quotation practice.

2.1 Quotations are names.

This view, attributed to Tarski (1933) and Quine (1940), considers quotations names. To make an example, if we consider the sentence ‘‘three’ has five letters’, we could replace ‘three’ by stipulation (we can replace it with a and then go on to say

that a has five letters⁵). The internal structure of the quotation is not considered important, and in fact the quotation is taken to consist of a simple expression. While surely this is the first step towards explaining **opacity** or lack of substitutivity (we cannot replace ‘three’ with a synonymous expression, if it existed, *salva veritate*), there are problems that militate against this theory. This theory obscures the grammatical and recursive operation through which quotation works (Davidson 1984). Furthermore, it does not fit in well with those cases of quotation in which we quote from an unknown language (in this case we are not sure what letters to use to express certain sounds). It does not fit in well with the one of the functions of quotation, which is to invent new words or coin new definitions, which are possibly infinite.

2.2 Quotations are descriptions. On this view, by Tarski (1933) and Quine (1960), a quotation is a description obtained by concatenating each successive letter of the materials quoted. One of the disadvantages of this is that we cannot quote from unknown languages (presumably because we are not sufficiently familiar either with their phonetics or written symbols) and that the quotations are largely unstructured. Thus this view does not do justice to the systematicity of language. (Surely an utterance is not only a concatenation of words but has a grammatical structure deriving from more general principles that are connected in a systematic way. This has been well known since the work by Chomsky; concatenations of words are not the only elements involved in quotation. When we quote a syntactically well-formed expression, we usually want to show what a person said and what he said usually includes the intention of forming a well-formed syntactic object). A further problem is that in many cases, quotations work both at the level of use and of mention, for such cases, this is not a suitable theory (John said that ‘Mary is crazy’(see also example (4)).

To amend some of the problems of this theory, Geach (1950) proposes that we should consider the quoted materials as concatenations of words. However, even if some structure is being introduced, this will not be enough, since we also need to

⁵ Someone took issue with this. However, since I am not going to defend the view of quotation as names, I will not reply to this objection.

know more about syntactic structure ('John went to the cinema' is a sentence of English).

2.3 The identity theory of quotation.

It has sometimes been claimed that in quotations, expressions refer to themselves or are identical with themselves. This view can be credited to Frege (1892) but also (exactly one century later) to Washington (1992). While this theory can in principle explain the asymmetry between written and spoken language, it is probably biased towards oral language. It considers quotational devices as not being indispensable. Nevertheless, the quotation operation is determined by a semantic rule: identify the expression with itself. This view has evident problems since, after all, we have multiple ambiguities showing up in quotation. Sometimes the expression is a phonetic form, sometimes a phonological form, sometimes a syntactic form, sometimes a lexeme or a string of lexemes, etc. (a point denied by Cappelen & Lepore 1999 and 2007, section 7.1.3, but pressed by Capone 2009a and Saka 1998, 123; 1999; 2005, 191; and 2011, section 6). So 'identity with the expression' does not really shed light on the multiple uses of quotation and it is clear that quotation does not consist of a simple rule such as 'establish identity with the expression'. (See also Reimer 1996, who explicitly expresses her reservations concerning the identity theory based on her idea that an expression does not really refer to itself, but in most cases of quotation refers to a type and in many cases one must establish what type is involved through pragmatic resources that mobilize contextual clues).

Another scholar who opts for the identity theory of quotation is Garcia-Carpintero (1994). This scholar considers the possibility of a pragmatic theory of quotation, but then he opts for the identity theory. I would like to answer the pessimistic view by Garcia-Carpintero on conversational explicature (actually he uses the term 'conversational implicature'). Garcia-Carpintero is aware that in various cases we could replace or supplement the demonstrative theory of quotation through pragmatic considerations. He considers a vocal utterance of 'Barcelona has 9 letters'. Since accepting that the city Barcelona has 9 letters will commit us to a falsity, we interpret the sentence as meaning that the expression Barcelona has 9 letters.

Carpintero's worry, however, is that once we accept a pragmatic view totally unconstrained by semantics, we are forced by theoretical parsimony to accept the identity theory of quotation. I would resist this theoretical step. The identity theory of quotation cannot work *unless* we supplement it with a pragmatic theory, and if we supplement it with a pragmatic theory, we shall have both a semantic principle and a pragmatic theory – which is less parsimonious than if we only had a pragmatic principle. The identity theory cannot work on its own due to the multiple (interpretative) ambiguities of the expression quoted.

2.4 Demonstrative theories.

The theory of quotation by Davidson is called the 'demonstrative theory' (the demonstrative theory has been renounced by its most prominent advocates, C&L). Davidson takes the quotation marks as the referring expression. The thing between quotation marks is the expression being mentioned or demonstrated. Davidson's theory is based on both the notion of demonstration and on demonstratives. An expression such as 'Cat' has three letters' comes out as 'The expression of which this is a token has three letters. Cat'. One of the original merits of Davidson's analysis is to have separated the quoted expression from the sentence in which quotation occurs. This is presumably useful in coining new definitions. (In coining new definitions, we need to keep the *definiens* and the *definiendum* separate; in a definition we invest a word with a meaning (or with a more exact meaning); when we say we shall use the term 'intentions' to refer to what the speaker intends to communicate, we are taking a word (or a term) and we are attributing to it a circumscribed meaning: thus it is convenient to separate the word we want to explain from the explanation, which is what happens when we use the demonstrative 'this' in the explanation and the term to be explained within quotation marks). There are, nevertheless, some shortcomings in this type of theory. It is not clear how one can account for recursive quotation as 'the'⁶. Cappelen & Lepore 1999 claim that one

⁶ One way to resolve Davidson's problem with respect to iterated quotation is tentatively something like the following, in connection with the example: 'the'. One could embed a second occurrence of a demonstrative through syntactic embedding in virtue of a relative clause and add subscripts to the various distinct occurrences of the demonstrative 'this': The expression of which this₁ is a token

should do without recursive quotation. However, it is clear that there is an asymmetry between the written language, in which recursive quotation clearly has a point and the oral language, where it is not obviously clear what its function would be. Given the fact that no quotation marks are actually used in the oral language, it is not clear how recursiveness can obtain. (It would also be difficult to keep in mind long iterated quotations in the oral language). (However, recursiveness should be explained with reference to the written language) Another problem of the demonstrative theory is that the quotation marks actually work as a subject. Thus, it would be possible in principle to have a sentence such as (3) in the written language:

(3) ‘ ’ has three letters.

But this is an impossibility. I suppose Davidson has to bite the bullet and say that (2) despite appearances is well-formed, albeit false. (What about ‘ ’ has zero letters’? For more on empty quotation, see Sorensen 2008 and Saka 2011.)

2.5 Mixed quotation.

A problem of great interest, raised by Davidson and taken up by Cappelen & Lepore (1997b), concerns ‘mixed quotation. It is clear that sentences such as (4) are meaningful and grammatical:

(4) John said that Mary is ‘a lioness’.

Cappelen and Lepore dismiss the view according to which one can append to the indirect quotation a phrase like ‘and the words which John really used are is ‘a lioness’’:

John said that Mary is a lioness;

And by the way, he used the words ‘a lioness’ in uttering what he said.

containing the expression of which this² is a token). (See also Gomez-Torrente 2001, 133).

Instead, Cappelen & Lepore capitalize on the fact that mixed quotation shows features of both direct and indirect quotation. Thus, what they do is combine Davidson's analysis for indirect and direct quotation. (However, notice that the analysis based on the parenthetical 'And by the way...' was dismissed because it multiplied entities. It is not clear that the double Davidsonian analysis is more reductive than the parenthetical analysis). Basically, in simple terms, Cappelen & Lepore's analysis says that (4) can be elucidated in terms of (5):

(5) There is an utterance *u* such that the speaker same-said the content of *u* and that is *u* and there is an utterance *u'* such that the speaker same-tokened the utterance *u'* and that is *u'*.

As far as I can understand, the virtues of the Davidsonian approach is that semantics and pragmatics are impossible to untangle, since the hearer must recover the referent of 'that'. While things might be easier in the written language, this would surely be a rather complex pragmatic task in the oral language. An approach along the lines of Geurts & Maier (2005) might be preferable. Geurts & Maier resort to pragmatic intrusion and account for cases of mixed quotation by accepting that there is more structure than is visible. Thus an utterance such as 'When in Santa Cruz, John orders '[ei]pricots' at the local market' is analyzed as 'When in Santa Cruz, John orders what he calls '[ei]pricots' at the local market'. This method cannot be directly applied to all cases of open quotation. See what happens in 'John said that Mary 'is a lioness''. I propose we can make use of an apposition structure, as in 'John said of Mary that she is a lioness, which he said by uttering the words 'a lioness''; anyway, a free enrichment view of mixed quotation looks preferable in the light of the fact that the convoluted double layer meaning à la Cappelen & Lepore is avoided. One could propose that we can paraphrase 'John said that Mary 'is a lioness'' with 'John said that Mary is what he calls 'a lioness''. These considerations, however, presuppose disambiguation, since a possible interpretation of this paraphrase is that John said that Mary is a lioness (the lioness is called 'Mary').

The explicature account of open quotation along the lines of Geurts and Maier appeared substantially flawed to an author like Saka (2011). I think Saka's considerations are not implausible, but they can be seen in a new light, taking into account Douven (2010) on the pragmatics of belief. Basically, Saka doubts that an utterance such as 'When in Santa Cruz, John orders '[ei]pricots' at the local market'

can be truth-conditionally equivalent to ‘When in Santa Cruz, John orders what he calls ‘[ei]pricots’ at the local market’. The reason for that is that what John calls ‘eipricots’ in Santa Cruz is equivalent to what he calls ‘aepricots’ in a different place. Geurts’ and Maier’s insertion of the materials ‘what he calls’ serves to create transparency, instead of opacity. However, the result of substituting coreferential expressions can be avoided if one follows, as is reasonable, Igor Douven’s (2010)

Epistemic Hygienics (EH): Do not accept sentences that could mislead your future selves.

It is possible to argue that Igor Douven’s Epistemic Hygienics, which is a pretty plausible principle, may be related to Sperber & Wilson’s principle of Relevance. Since, after all, the explicature through which pragmatic intrusion is constructed is a pragmatic process, the Epistemic Hygienics is clearly relevant and applicable and it furnishes the right results in the case considered. In fact, by replacing ‘[ei]pricots’ with ‘[ae]pricots’ in an explicature, one obtains a result which is likely to mislead one’s future self. Replacing ‘[ei]pricots’ with ‘[ae]pricots’ has the immediate result that in the future the speaker can no longer remember which word(s) John used when buying apricots in the local market in Santa Cruz.

2.6 Decisive objection to demonstrative theories.

We have already seen the shortcomings of demonstrative theories of quotation. The recursion problem noted and dismissed by Cappelen & Lepore is a serious shortcoming. But the most thorny shortcoming of all is that demonstratives and demonstrations work well in the case of the written language, where we have quotation marks and inscriptions between quotation marks, but much less well in the case of the oral language. Demonstrations and demonstratives require pointers but it is not clear that spoken words can be used in this way. First of all, because of the lack of pointers. Second, because it is not clear what and where the pointers should point to. If this was not enough, in the same way in which the identity theory of quotation is jeopardized by the ambiguous uses of quotation, the demonstrative theory has a parallel problem. Even if pointers were available, and one could point towards a location, it is not clear that semantics is responsible for identifying the metalinguistic level required by this particular quotation. Even if we accepted the

machinery of Davidson's ingenious semantic treatment, we would have to amplify it by means of pragmatics. Pragmatics would have to do three things in the oral language:

Provide pointers;

Point at something physical;

In virtue of pointing at something physical, point to a particular linguistic object.⁷ (This point recursively involves pragmatics as one needs to explain how one can refer to a type by referring to a token).

⁷ Nunberg (1993) says that whenever one points to something, a distinction must be made between the thing which is pointed to and that which is referred to through the pointing. Recanati (2010) takes advantage of this distinction to show that by quotation a speaker can, in virtue of demonstrating a certain token, point to a certain type (an expression type). However, it is clear that adoption and exploitation of Nunberg's distinction simply serves to add levels to the analysis of quotation. An analysis in which demonstrating something is not indispensable must be preferred, since invariably the notion of demonstration requires the meaning transfer implicit in Nunberg's analysis and a shift from a token to a type in virtue of some pragmatic mechanism. An analysis that gets rid of these two steps is surely preferable. Reimer (1996, 139) who argues that the demonstrative account is empty, like Nunberg argues that contextual clues help effect a transition from a token to a type:

There is no need, for instance, to actually point to a particular copy of *Word and Object* when I come out with my utterance of (4) [That was published in 1960] – if, for instance, it is the only book on the bookshelf that my audience and I are viewing. The context, even without ostension, is sufficient to 'demonstrate' (to display) the book in question. And it seems quite plausible to suppose that, in cases of spoken quotation, something similar is going on: features of the context other than demonstration effectively 'demonstrate' the uttered token, the latter of which is used, in a process of deferred reference, to refer to the corresponding utterance type.

I think that the flaws in this type of analysis are evident. Due to the discrepancy between written and oral language, Reimer is compelled to say that it is not the demonstration that does the job of demonstrating. Nevertheless, she accepts some mysterious notion of demonstrating something in virtue of contextual clues, and this something is a token. Her analysis, as is clearly visible, involves at least an unnecessary step: reference to a token. My pragmatic analysis, instead, would dispense with the notion of referring to a token in order to refer to a type, and immediately provide the level of meaning required through explicature and the contextual resources which explicature mobilizes. In particular, there is no need to use the logical step of referring to a token in order to refer to a type, in the same way in which, in the case of metaphorical expressions, nobody (in pragmatics circles) really accepts that one moves from a level of literal meaning to a level of intended meaning.

It is clear from this analysis that the pragmatic tasks have been unnecessarily multiplied.

The pragmatics of quotation can avoid the first two steps. The third step also needs drastic reformulation, since one needs to get rid of the notion of ‘pointing at’, which seems to require the use of demonstratives. All we need is to recover the speaker’s intentions and this can be done by resorting to the multiple clues and cues used by the speaker

So, in a sense I am not only arguing against theories of quotation that make use of demonstratives, but also against Recanati’s theory of quotation, since this involves demonstrations that exhibit a certain type. My view is that singling out the metalinguistic type involved in a particular quotation does not require pointing or demonstrating. Functional applications of a predicate are much more efficient than pointing, since it is the fact that a certain predicate applies to a certain type that gives us for free the thing which, for other theorists, is being pointed at.

2.7 Recanati’s theory of quotation: towards a pure pragmatic view of quotation.

At this point I can very well conclude my short summary of theories of quotation by briefly discussing Recanati’s ideas. Recanati’s theory is clearly most indebted to pragmatics. He still uses terms like ‘demonstrating’ but, in fact, he abandons the Davidsonian theory by giving it a more pragmatic dimension.

The key point in Recanati’s analysis is that an expression is being demonstrated by being quoted through the device of *conventional implicature*. But this view is immediately tainted by the ambiguity problem. What is being demonstrated (assuming that it is being demonstrated) is not just an expression, but what Bennett (1988) calls a feature associated with an expression. A conventional implicature analysis would work well if one could show that the thing demonstrated is an expression. But this is not the case in many cases. Sometimes it is just a phonetic string, which is not an expression in the full sense of a linguistic string having a form and an extension. Sometimes it is a lexeme. Sometimes it is a sentence. The

things capable of being ‘demonstrated’ are too disparate to allow one to hope that they can be unified. And if they cannot be unified, it cannot make sense to say that a quotation singles out an expression by conventional implicature. The conventional implicature analysis would necessarily require complementing by a conversational implicature analysis. And this will immediately lead to unnecessary multiplication of levels. Furthermore, one problem with this view is that conventional implicatures, according to Grice at least, require a linguistic trigger. Words such as ‘but’ and ‘therefore’ are cases of triggers of conventional implicatures. But these are words associated by conventions with certain meanings (even if they operate at a level of procedural meaning). It is not clear which words would have to carry the same semantico-pragmatic task in the case of the conventional implicature involved (somehow) in quotation. While Davidson’s analysis has an element which carries out the task of selecting the quoted expression, namely punctuation marks that function as a demonstrative, there is no equivalent word in Recanati’s account. The fact that quotation marks are used in the written language surely does not prove much, since they could be replaced by italics⁸, and then in this case it is not clear which elements would convey the alleged conventional implicature. But the most crucial problem is that, in the oral language, there are no quotation marks at all; in fact, it is even possible not to have quotation marks in the written language. So, it is not clear that there is a suitable candidate for a suitable trigger at the level of words for the operation of quotation.

Wayne Davis (personal communication) has one objection to my view: “I did not understand the problem you see with the ambiguity of quotation. Whenever anyone uses a demonstrative pronoun like ‘this’, there are many possible referents compatible with the meaning of ‘this’ even when supplemented with a pointing. So why is the fact that there are several possible referents in any case a problem with the claim that quotation involves demonstrating an expression?” My answer is

⁸ Why does a conventional implicature have to be conveyed by a word rather than a feature like italics? Suppose that italics is used to express the conventional implicature associated with quotation. Then suppose I say or write: This is what John wrote today: *I am not happy*. Then italics would have to be associated with quotation, but in this case italics could not quote an italicized item. And in my utterance the purpose of my quotation was to quote what John wrote, including his use of italics. I do not see how the conventional implicature analysis can solve this problem. The pragmatic hypothesis faces no such problem, since italics could just be used to show how John wrote the utterance, by citing the font he used. That the teacher intends to quote the font as well could be fixed by *intentions*.

simple. Demonstration is canonically used in contexts where, by demonstrating an object or an individual, one provides an answer to a question (whether voiced or not). So, if you ask me, who is John, and I answer ‘That is John’ (only John and Mary are visible), I use the demonstration informatively, because my act of demonstrating an object has unambiguously picked up the right individual. In case there are a hundred boys concentrated in the same courtyard, and I say ‘That is John’, pointing my finger at the group, there are not many chances that my demonstration act will successfully demonstrate anything. I will probably have to call: ‘John!’ and when he replies, I could then demonstrate ‘That is John’. But at this point the demonstration is quite superfluous. My argument for quotation is the same. Demonstrating is superfluous and we need alternative pragmatic means for establishing what the quoted item is. (In support of the conventionality of quotation, however, one could say that whether or not demonstration is superfluous, quotation marks succeed in establishing what the quoted item is. But I am not too sure that this is right. Take the following case. ‘This is what John said: ‘ ’. Which is the quoted item? Could it be John’s silence? Could it be the lack of an item to be quoted? This is a terrain where pragmatics has certain advantages).

3. Towards a pragmatic view of quotation

Furthermore, a pragmatic theory which accepts that quotation is done through conventional implicature, but must admit the shortcomings of this hypothesis, by remedying them through a reparative conversational implicature is surely not on the right track, since it comes short of parsimony and Modified Occam’s Razor. (The pragmatic view alone suffices).

We need a level of implicature, also to exclude unintentional features of the reported message, as Bennett (1988) says. Consider the written language only: if Sally writes (6)

(6) John wrote: MARY must NOT come

then she will probably implicate that there is a reason why she is reporting some of the words in boldface and the reason must be that this is the way the original message was printed. Following Sperber & Wilson (1986), the cognitive effort involved in forcing the reader to distinguish boldface must be offset by appropriate cognitive rewards, and the easiest way to make sense of the additional efforts to which the reader is put is to convey that the reporter wants the reader to know how the words were originally printed to start with.

Wayne Davis (personal communication) does not agree with my explanation. It is possible that the extra effort is balanced by the intention to convey the words highlighted were of particular importance to the reporter or that the reporter wants to highlight a contrast with what Jane said (Mary will come). Or perhaps the reporter is merely trying to highlight words that were nearly illegible when Mary wrote them. It is true that it is not impossible to have other interpretations than the one which immediately came to my mind. However, it should be possible to order interpretations according to their relevance. Since the font used in the report is of immediate relevance to the question of how John wrote the sentence, I think this intention has priority over the others, which probably require a greater amount of contextual clues to reinforce them (or to promote them to intended interpretation).

Perhaps the most problematic point of pragmatic theories is the controversial claim that (5) displays English words, that the words were not originally printed in Latin. However, once we abandon the tenet that quotation is a semantical device and that it always consists of quotation marks (whether explicit or implicit) that point to tokens, we can make sense of this apparently bold pragmatic claim. The fact that a quoted text may have originated in a different language is a challenge to semantic theories of quotation because it shows that the objects which quotation may point to may be multiple and can be more abstract than the tokens quoted, thus requiring principles of pragmatic interpretation, rather than semantic rules. In order to explain why the canonical interpretation of quotation is to select a fragment of text in the same language in which the reporter utters it, we need to invoke Relevance Theory again with its insistence on Relevance as the balance of cognitive effects and processing efforts. Since the reporter did not quote the words uttered in a different language, the cognitive effects are large enough if the linguistic expression

originally used is English. Placing greater cognitive efforts on the hearer by using a different language, would have to be offset by greater contextual effects.

A problem which Recanati faces in responding to considerations about mixed quotation in Recanati (2008) is that, in fact, the implicatures associated with mixed quotations are not cancellable. Consider (7):

(7) John said that Mary is ‘a bore’.

If one tries to cancel the interpretation that John uttered the words ‘a bore’, one fails to do so, according to Cappelen & Lepore (2005b). One could try to rescue Recanati’s pragmatic analysis by saying that, after all, ‘a bore’ need not have been uttered by John. It is possible that the reporter was echoing someone else and that the reporter is stressing this echoic relationship. This is not implausible. However, at the end of a chain of inference, we will end up with someone who actually uttered the words ‘is a bore’ and we will also have to admit that there is an echoic relationship between Mary’s words and that echoed person⁹. But this in practice means that Cappelen & Lepore are right in insisting that the explicature (as the level of meaning resorted to in quotation is an explicature) cannot be cancelled. But this is not something undesirable¹⁰. These data fit in well with the picture of non-cancellable explicatures which I traced in Capone (2009a). What I pointed out in that paper, in its essentials, is fundamentally that the speakers’ intentions, once they are manifested, can no longer be cancelled. The reason for this is that a speaker uses many contextual clues in getting an intention across, and once these clues are *disseminated* or made available through a text, it is impossible to go back and

⁹ Saka (2005, section 3.3) rejects the echoic account on the grounds that it does not strike him as plausible that the word choice signaled by (echoic) quotation must reflect someone’s thoughts in order for the quotation marks to make sense. The example Saka addresses is quite cogent. If John says ‘The ‘debate’ resulted in three cracked heads’, there need not be someone who described the event as a ‘debate’, but John may plausibly mean that the event in question started as a debate and then degenerated into a violent discussion. I agree with Saka that speaker’s intentions are even more pervasive than Recanati allows. However, my point remains unchanged: speaker’s intentions, once manifested, cannot be abrogated, resulting in an explicature that cannot be cancelled.

¹⁰ Recanati’s reply to Cappelen & Lepore is that ‘Cancellability is a *necessary* condition for a meaning component to count as pragmatic; it is not a sufficient condition.’ So cancellability *per se* does not establish the pragmatic nature of the meaning component’ (Recanati 2010, 208).

withdraw an intention. The intention can be retracted, but never cancelled. I have shown in my paper on the attributive/referential distinction that referential or attributive interpretations cannot be cancelled; and this is due to the richness of the clues disseminated in a text. Since intentions can only be fixed pragmatically, literal meaning being essentially inadequate in fixing intentions without the help of pragmatics, what was manifested as intended cannot be undone¹¹. It can be retracted, but this does not mean that an explicature disappears. Consider the following example from quotation:

(8) 'Red' has three letters

In interpreting (8), what we clearly do NOT do is to establish an identity between an expression and the expression quoted, because there is not such a simple rule as this and we have already said that the word 'expression' is ambiguous. By 'red' we do not pick out a word, but we pick out a written form. In other words, we need to *subtract* the meaning level from the lexeme before arriving at the intended linguistic unit. First of all, we establish that the unit 'red' is semantically inert, in other words that we are focusing on something different from its extension. But we do so not through Washington's identity rule, but through pragmatics. We exclude that the extension of 'red' is being intended because the statement 'red has three letters', interpreted as being one about the red color, is blatantly false. Once we exclude that a referential entity is intended, we need to find out which feature of the expression is being focused on. Since 'three letters' immediately activates encyclopedic knowledge (that the things which can have letters are written strings), we will select a written string. Notice that in the complex cases succinctly described through pragmatics, we do not demonstrate anything at all, nor do we use implicit demonstratives. This does not amount to excluding demonstratives from playing a role in quotation, as one can clearly have things like:

(9) Look at the way this is written,

¹¹ Wayne Davis (personal communication) argues that in the case of 'John said that Mary is 'a bore'' the implicature that John used the words 'a bore' is cancellable. In fact, he may have used the Italian equivalent of those words. But look at what happens if we know that John does not speak Italian or any other foreign language. In this case, it becomes very difficult to cancel the explicature.

where one demonstrates an inscription or a shape by pointing at a sample of writing . Of course this possibility exists, but it is more naturally confined to the written language and to particular types of inscriptions.

The things I said about the non-cancellability of explicatures seem to reinforce the picture of pragmatics and modularity which I presented in Capone (2010c) and in Capone (2011). In Capone (2011) I replied to Cummings (2009) and I said that pragmatic processes of the non-reflective type are non-cancellable, are driven by fast-and-frugal heuristics and by mandatory operations which are bounded by the principle of Relevance, which operates by throwing a net on information that can be processed and which could direct or orient the pragmatic interpretation process. Of course, I granted that parts of pragmatic processes need not be mandatory or encapsulated, especially if we focus on reflective inferences. I doubt that reflective inferences are involved in quotation. It appears that inferences involved in quotation are fast, mandatory, automatic and encapsulated. This is the way the mind works in whatever language. We could make an experiment to test for encapsulation. Suppose you stipulate that quotation marks, in a certain fragment of the language, mean ‘Start looking for an expression which is located two meters away’. Present the subject to be tested with an example such as ‘‘Cat’ has three letters’. Then you need to guess whether quotation will be interpreted along cognitive routines dictated by universal cognitive principles or whether quotation marks will be interpreted in a procedural way by following the explicit stipulation. My guess is that the interpretation of quotation follows a different path of interpretation from stipulated interpretative routines. Paolo Leonardi (personal communication) is not persuaded by this thought experiment. He says: “The experiment does not persuade me at all, because we do not see any expression two meters away and because to stipulate something out of the blue and find that it is unnatural, when there is a usage that is established by a linguistic “convention” is predictable whatever the experiment” (translation mine). There are two things which do not persuade me about Leonardi’s line of thought. First of all, if I see a certain bus at a distance and I see that it is the bus number 8, I have seen an expression at a distance (of course I have seen a token of that expression; but it is not impossible for me to say that expression is a number, in which case I am not pointing to the token, but to the type). Second, Leonardi assumes that the meaning of quotation is provided by a convention (a linguistic convention, in fact), which explains why the experiment of stipulating a certain

meaning for the quotation marks is not a happy one; the conventional meaning is always to be preferred to the arbitrarily stipulated meaning. However, if one does not assume that quotation is regulated by a linguistic convention (as I do), the experiment makes sense. Let us accept the experiment only conditionally, as indicated by Leonardi. If the view that quotation is expressed by pragmatic means is accepted, then it will follow that the pragmatic inference will be mandatory.

I anticipate the objection that I have not (yet) said how quotation is to be interpreted in written cases. My answer is simple. Quotation works like all cases of marked expressions, where the use of a marked expression signals a deviation from an ordinary pattern of usage. (It is tempting to see analogies with cases of anaphora by Levinson (2000) and Huang (2007)) Cognitive efforts must be offset by cognitive rewards, if we accept the theory by Sperber & Wilson. Thus one way to make sense of the marked pattern is to guess that it correlates pragmatically with an interpretation that is alternative to the interpretation that normally occurs in ordinary uses – an alternative to the extensional use of an expression. This alternative obviously has a metalinguistic character.

Of course, this is not the end of the story. I quite agree with Wayne Davis (personal communication) that, if this was the end of the story, then we would have to account for why in, e.g. ‘‘red’ has three letters’ we do not choose an interpretation according to which the sentence means that any word with the same meaning as the word quoted has three letters, or that the word in quotes read backwards (‘der’) has three letters, etc. It is clear that pragmatic considerations of the contextual sort play a role here. In particular, encyclopedic knowledge plays a role in enriching the explicature.

Could we explain quotation by resorting to default semantics (Jaszolt 2005) ? Here I am not sure that I have a definitive answer. I have argued in Capone (2011) that items of the lexicon that are associated with certain default pragmatic interpretations are stored in a default semantics archive. However, in this case, we have no word since the quotation marks are interpreted (see Saka 1998 and Washington 1992) as punctuation marks. They exhibit structure, and they are treated on a par with subscripts used in syntactic notation to signal phrase markers. It is therefore difficult to assimilate quotation marks to items that can be characterized through default semantics. Although I do not admit many contextual elements in my

view of quotation, I predict the interpretation of quotation entirely through pragmatic principles. I assume a certain systematicity in the way pragmatic effects are obtained.

3.1 Quotation and cancellability

Now, before getting into the details of quotation, I want to establish the particular notions by reference to which I will discuss my views of quotation. For years, I have voiced an unorthodox view of explicatures saying that this view has the immediate advantage of solving Grice's circle¹², a problem noted by Levinson (2000). In short this view is that explicatures are not cancellable. I have usually defended the weaker version of this view, distinguishing between potential and actual explicatures, and saying that only actual explicatures are not cancellable. While I will not rehearse all the arguments used to persuade readers of Capone (2009a) on non-cancellable explicatures, suffice it to say that explicatures, on my theoretical view, arise especially in those cases in which the semanticist is forced to say that literal meanings lead to contradiction or patent falsity. To rescue an apparently plausible and well-formed sentence from accusations of falsity or contradiction, one needs to posit an explicature. This view is more parsimonious than those that are around, in line with my idea that one should try to adhere to a classical semantic picture whenever it is possible, preferably by using the armory of implicit indexicals or variables. To keep this story short, I believe that, in general, semantics provides the guidelines for the use of a linguistic expression and tells us or instructs us on how to use it and how to interpret it. I believe that sentential meanings are not necessarily gappy (that is, they cannot serve to express a truth-evaluable proposition), even if sometimes they happen to be gappy, in which case it is our task to show that pragmatics is necessary to explain the phenomenon of pragmatic intrusion, how what is said is said despite a shortfall of the semantics of the sentence.

I have tried to connect my ideas about the non-cancellability of explicatures with views on modularity of mind and their intersections with the pragmatics of language. Unlike some orthodox views in pragmatics (e.g. Kasher 1991 or Cummings 2009), but like other eminent pragmatics scholars (Sperber & Wilson 2002, Wilson 2005,

¹² Conversational implicatures take input from truth-conditional meaning but truth-conditional meaning depends on pragmatic intrusion; thus, we have a double application of pragmatic principles.

Carston 1996) I accept that at least part of pragmatic interpretation is modular in nature and I respond to the standard objections to the modularity of pragmatics, by expatiating on the virtues of my view according to which explicatures are not cancellable. I showed elsewhere that fast-and-frugal pragmatic heuristics determine interpretation of a linguistic expression on the spot, without the necessity of re-opening the interpretation case whenever a piece of information that is potentially in conflict presents itself. The reason for this is that interpretation processes are bounded by the Principle of Relevance, which is like a net-throwing principle (it filters away information that does not interact in a fruitful way with assumptions being considered) and determines the availability only of information relevant to the interpretation process. The fact that in my view explicatures are not cancellable, seems to conflict with the idea by Kasher that pragmatic inferences are not encapsulated because their defeasible nature means that every piece of information drawn from encyclopedias can interact with existing interpretations leading to re-interpretations (or modified interpretations). Non-cancellability and encapsulation seem to go hand-in-hand on my view. It may be thought that my views on pragmatic intrusion are in conflict with Recanati's view about truth-conditional pragmatics; however, there are various points in Recanati (2010) which seem to indicate that pragmatic intrusion cannot be easily cancelled; he is at pains in explaining these cases away, while if one were to accept my view that intrusive explicatures cannot be cancelled, one would not be at pains in explaining the points that are in conflict with the assumption that pragmatic increments are cancellable. My views about pragmatics are best exemplified in Capone (2009a), where I did my best to explain that a theory of pragmatics should be a theory of (communicative) intentions and that intentions are fixed in context through a number of clues and cues disseminated by the speaker and utilized by the hearer. Once intentions are fixed and messaged are implicated, they cannot be un-implicated (as Burton-Roberts (personal communication) would say).

In this section I will show why my considerations on the pragmatics of quotation important from the point of view of the general theory of pragmatics. I want to show that understanding the issue of quotation will have a bearing on the way we understand pragmatics and, in particular, the issue of cancellability of explicatures.

I will keep this section short on purpose, because, otherwise, I run the risk of developing a separate paper. However, this section contains all the essentials of my proposal. I mainly want to address a reply by Carston (2010) to Burton-Roberts (2005) on non-cancellability of explicatures. Since my treatment of explicatures and my discussion of cancellability intersect with considerations by Burton-Roberts, I take Carston's objections as applying, *mutatis mutandis*, to my own views. First of all, Burton-Roberts considers the possibility of cancelling an explicature a logical impossibility. Why should this be so? Let us remember that the need for an explicature is dictated by the fact that a sentential meaning is often gappy. Thus, we predict, on logical assumptions, that cancelling an explicature from the utterance of a sentence whose meaning was originally gappy (to start with) cannot result in contradiction. This is a process quite unlike implicature cancellation, which usually results in an utterance which is not apparently contradictory, but where the possibility of contradiction is real. There is no way to disprove the theory, if we start with the assumption that a sentential meaning is gappy or underdetermined¹³. Carston opposes this sophisticated reasoning, by saying that her students normally cancel implicatures/explicatures as in 'She's ready, but Karen isn't ready to leave for the airport' and find the following 'She's ready but she isn't ready' contradictory, despite the fact that the sentential meaning is gappy. So, Carston's point is that people can and do, quite confidently, assess examples like the ones just given for contradictoriness and can express judgements of non-contradictoriness or contradictoriness.

Readers, at this point, will realize that it is not easy to make a theoretical choice. We are at an impasse: theoretical considerations vs. intuitive judgments by normal speakers. I propose to suspend our judgments for the time being, even because Burton-Roberts claims that in many cases explicatures cannot be cancelled without conveying a sense of contradiction (as Carston says, his position is even stronger, since he also takes cases of particularized implicatures not to be cancellable, on the

¹³ In other words, contradiction is a property of statements, not of sentences and thus the cancellability test cannot be applied to sentences; an explicature is usually added to a sentence to transform it into a full proposition; but then we cannot apply cancellability to return to the original sentence. To apply cancellability, we would have to show that the sentence and the actual explicatures are not contradictory; but to show they are contradictory we need to compare statements, not a sentence and a proposition (the explicature)). A priori, we predict that cancelling the pragmatically derived elements of an explicature will not result in contradiction.

grounds that once intentions have been fixed, they cannot be revoked). At this point, Carston admits that there are problems. However, she insists that cancellability for Grice was only a diagnostic for the notion of conversational implicature (and not even the most important one, given the decisive role played by calculability) and was based on two main ideas: conversational implicatures can be either explicitly cancelled or contextually cancelled. Burton-Roberts has only shown that explicatures cannot be explicitly cancelled. (Anyway, he has pointed out cases in which it is difficult to cancel an explicature). However, he has not demonstrated that explicatures cannot be contextually cancelled.

Things are not that easy. For instance in Capone (2003b, 2006 and 2009a) I have argued that explicatures cannot be cancelled on various theoretical grounds. Part of the evidence comes from accepting that explicatures are often a tool for resolving a theoretical problem and for showing that an utterance which otherwise has the appearance of a falsehood or contradiction, is not false or contradictory. If the explicature is a tool for rescuing a particular utterance by resolving a potential contradiction or logical absurdity, it is clear that cancelling the explicature turns out to be very problematic.

Carston, however, insists that Burton-Roberts and I are misguided as we consider utterance processing to be an online process, while the notion of cancellability has more to do with its being a (theoretical) diagnostic, a test that must be passed by simply considering potential rather than actual explicatures.

Now, in my 2009a paper I was prepared to grant that one has to make a distinction between potential and actual explicatures and that my notion of non-cancellability of explicatures was limited to actual explicatures and, presumably, did not affect potential explicatures. However, my current perspective on quotation compels me to consider or reconsider whether what Carston says about the cancellability of explicatures on the basis of contextual cancellation has full generality. And my reply is not positive. Consider what happens when we have explicatures relating to vocal utterances such as 'Bold has four letters'. Given what Carston says, it would be possible to cancel this explicature if we changed the context (this is what contextual cancellation amounts to). But it is not clear how we should change the context in this case. Presumably, we need a context in which our assumptions of normality are suspended. Consider John, the madman. Is he so mad to be able to attribute the

predicate ‘has four letters’ to the person who is bold (assuming an interpretation of the somehow defective utterance ‘The bold has four letters’ which we may be inclined to attribute to him)? Seeking a context that legitimates the cancellation of the explicature leads us to a situation where we have to suspend judgments of rationality and if we have to suspend such judgments, then we probably had better give up our attempts to build a theory of rational interpretation of communication. After all, our theoretical assumptions are based on rationality (even Relevance Theory, which is less reflective than Grice’s pragmatics involves principles of rationality such as judging the amounts of cognitive efforts as compared to the amount of cognitive effects). So, such a drastic departure from rationality will not do our theory any good. In my view, searching for contexts in which quotational readings can be cancelled is not easy. As I said, even in the cases of open quotation, where a portion of the utterance is cited (hence attributed to another person) one mobilizes resources that lead the hearer to understand the explicature and, thus, undoing the explicature is irrational in that it will annul the systematic and intended effects of mobilizing resources, a process which has both a production and a comprehension cognitive cost. The very idea that cognitive costs are decisive in interpreting utterances and in determining explicatures runs counter to the theoretical option of having cancellable explicatures. This is an extremely costly move, as Jaszczolt (2005) admits.

3.2 Compatibility with Saka’s view.

Before closing my discussion of explicatures, it will be important to consider if and to what extent my treatment of explicatures is compatible with Saka’s view of quotation. Saka (1998) may well accept my explicature account (or part of my explicature account), but insist that a minimal semantics for quotation must be provided, and that is what is supplied by the identity principle and its rule that an expression is identical with its nonce referent. Saka may very well insist that we can minimally differentiate use and mention and that mention involves identity of an expression with its nonce referent. Saka minimally characterizes use and mention in the following way:

(u)

Speaker S uses an expression X iff:

- (i) S exhibits a token of X;
- (ii) (ii) S thereby ostends the multiple items associated with X (including X's extension);
- (iii) S intends to direct the thoughts of the audience to the extension of X.

(m)

Speaker S mentions an expression X iff:

- (i) S exhibits a token of X;
- (ii) (ii) S thereby ostends the multiple items associated with X;
- (iii) S intends to direct the thoughts of the audience to some item associated with X other than its extension.

Now, while I do not doubt that this is a laudable definition of the distinction between mention and use, one that is interesting from a philosophical point of view, I doubt that this is a semantic distinction adopted by the speakers of a language like English or Italian. In fact, if, as I assume, the pragmatic process has as its aim the attainment of the most specific interpretation possible, it will be clear that on all possible uses of mentioned expressions a speaker commits herself to (m) in virtue of the logical implications of the explicature. Consider what happens when one hears the utterance ‘‘shhhh’ is a pleasant sound’. One clearly concentrates on the phonetic string and thus by explicature it can be assumed that (m) obtains. To have a pragmatic process that starts with (m) and then is enriched by the ulterior inference that we are talking about a phonetic string and not about the normal extension of a word involves an extra step. Furthermore, adopting (m) as a basic semantics involves knowing whether one is faced with the task of first disambiguating pragmatically whether rule u or m is applicable and then adding further pragmatic increments. One, in other words, needs at least two pragmatic steps, plus a semantic step, while a single pragmatic step is preferable. (Saka's work is not presented as a theory of interpretation.

However, in this chapter we face the task of providing a theory of interpretation based on the speaker's intentions).

3.3 Camouflaged direct quotations.

People occasionally say things like:

(10) I 'almost killed her'

Using some contextual clue to indicate that some portion of the utterance does not belong to them. One can, for example, use the device of mimicking somebody else's voice (as pointed out by Recanati 2010) to mark a segment of one's utterance as representing somebody else's voice or point of view.

In written cases, one can use orthographic devices to impute certain segments of one's utterance to somebody else (see the examples by Dickens quoted by Recanati (2010) which involve extensive mimicking).

If we want to opt for a unified treatment which handles these examples just mentioned in the same way as the cases of explicit direct quotations (to be more precise, where the verb 'say' or a similar verb is used), we need some way of marking off a segment from the remaining part of the utterance and conversationally implicate by stylistic marking that the segment is marked for special interpretation, a case of mention rather than a case of use. Stylistic marking is important and does not necessarily involve mimicking or closely imitating the model of the voice to be reproduced. All that is required is that one uses a different style. So, one could differentiate some serious portions of the utterance from some segments uttered in a style that is evidently, blatantly, intentionally different from the rest: say a joking tone of voice, laughter, etc. Stylistic marking is important and suffices to trigger an implicature through the heuristics we are used to. What is marked is not used in a normal way and carries messages concerning the point of the usage.

The gist of Recanati's treatment of ironic distancing from certain portions of the utterance or assuming an authoritative voice by imitating or representing an authoritative voice, is that one implicates the reason for the departure from ordinary usage (what he calls the 'quotational point'): in other words, one implicates what

intended effect one wants to achieve (or the purpose of the quotation). Consider again (10) (I ‘almost killed her’). How should we interpret this? One natural way to interpret this is to make use of Geurt’s treatment of mixed quotation and to allow for pragmatic intrusion. In other words, we could render an utterance of (10) as (11)

(11) I did something which she described/would describe as almost killing her.

I would like to call (10) a case of camouflaged quotation because the ascriber is hidden and because part of the intended meaning must be reconstructed somehow by inserting some materials. Clearly in (11) contextual considerations are responsible for someone’s opting for the past tense or the conditional. In this case, we can clearly see that there is a division of labor between default inferences (the triggering of a departure from normal usage through the usage of a stylistically marked portion of the utterance) and contextualized or particularized inferences. We need to know what the speaker has in mind, who uttered that segment of utterance in the past, and other relevant pieces of information allowing us to recover the purpose of the quotation and its intended source. In this case, the speaker probably wants to play down the gravity of what he has done by implicating that the source is in the habit of exaggerating things, of amplifying them in a disproportionate manner. I have no objection to Recanati’s wish to retain the source and the purpose of the quotation as particularized implicatures, rather than generalized ones, provided that it is made clear that the full interpretation of the quotation requires 1) a segment of the utterance which is sufficiently marked to produce M-implicatures (implicatures due to the use of a marked expression, an expression which deviates from the unmarked form in that it is more prolix or less frequently used); and 2) an explicature which reconstructs at least partially the speaker’s communicative intention in connection with the explicated articulation of the utterance. This would normally insert words such as ‘I did what X (the source) would be prepared to describe/described as Y’.

A remaining and not negligible question is whether it is possible in all cases of implicitly quoted sources to articulate the implicit portion of the utterance through an explicature along the lines of Geurts and Maier. Consider the utterance ‘Einstein arrived’. Suppose this utterance is prompted by a situation in which a person is under the delusion of being Einstein. Then this utterance could mean that the person who believes himself to be ‘Einstein’ arrived. This could be taken as the implicature/explicature of the utterance. The implicature/explicature of the utterance

would involve reconstruction of the agent through an objective description and then attribution of some mode of presentation of the referent which uses the words likely to be uttered by the source. It is a question of time and patience to see whether all examples can be dealt with in terms of explicatures; but it should be said that if things were this way, an important aspect of the theory would involve emphasis on the explicature which would partly be due to generalized implicatures and to contextualized inferences.

We may very well wonder if camouflaged quotations are cases of cancellable or, otherwise, non-cancellable explicatures. I have said that the explicature involves two stages: spotting a segment that is stylistically marked; attributing a source on the basis of contextual information. The two tasks are probably connected, depending on how closely the marked portion of the utterance mimics the source. We have no doubt, here, that there are enough contextual clues to set the interpretative enterprise on a non-return trip. If a segment of the utterance is stylistically marked, there must be a reason for that, which invariably triggers the implicature. If a portion of the utterance is labeled by common knowledge as belonging to a certain person, who is the source in that he always used those words in the past, then there are enough contextual clues to identify the source. The explicature thus obtained is clearly non-cancellable. Lack of cancellability does not have to do with conventionality (conventionality may play some role but not a decisive one), but with the fact that if one mobilizes the semiotic resources towards a pragmatic inference, it is difficult, if not impossible, to go back and pretend that those resources were never mobilized, that there was no reason in disseminating the text with abundant interpretative clues. The theory of relevance predicts that it would be uneconomical to cancel an explicature obtained by mobilizing contextual clues, because the cognitive efforts involved would have to be expended in vain in case an explicature has been cancelled.

3.4 A possible objection from Predelli (2003a).

Predelli (2003a) examines a limited amount of uses of quotation (mainly cases of mixed quotation) and stresses the analogy between quotation and conventional implicature. The structure of his paper follows closely the structure of the work by

Recanati on quotation: on the one hand, he assumes that quotation requires an analysis in terms of conventional implicature; on the other hand, he allows that quotation involves some elements of interpretation that are of a deep contextual nature (presumably the point of the quotation). What is really of interest in Predelli's proposal is that, in the same way in which conventional implicatures are cases of a doubled layered meaning, quotation too, for him, involves a double layered meaning structure. To make things clear from the start, consider (12):

(12) Mary is rich but generous.

This statement has a truth-conditional layer of meaning, which can be paraphrased as 'Mary is both rich and generous' in addition to having a secondary attached meaning, such as 'There is a contrast between being rich and being generous'. The two layers of meaning need not be of equal importance, and Predelli, in line with an authoritative tradition on conventional implicature, accepts that the truth-conditional content, but NOT the attached conventional implicature, can be the target of negation. Thus, if you say 'That is false!' in reply to (12), you are apparently negating that Mary is both rich and generous. (You are not denying that there is a contrast between being rich and being generous). I have no particular complaints concerning the validity of these considerations in connection with conventional implicature. However, I am not completely persuaded that these considerations apply to quotation fully. At this point, I have to note that the examples used by Predelli to support his ideas are admittedly limited, as he just uses some cases in which the quoted item is also used. Quotation, on the other hand, in its most general significance, may also include cases in which the quoted item is mentioned but not used. To start with, however, we can very well check whether the data considered by Predelli really support his view of quotation as conventional implicature. Consider:

(13) These are the 'proofs' of the illustrations.

Predelli considers a metalinguistic usage of 'proofs'. This usage is not correct (as it is not standard) and the speaker knows that well, and so it is marked as deviant through the use of quotation marks. (But of course, when the quotation marks are added, the utterance is not deviant at all, but somehow indicates that the utterance without quotation marks would be deviant). According to Predelli, if speaker B replies 'That's false', he cannot target the quotational use by negation, but only the

truth-conditional content of the utterance. Predelli must be wrong, because after all, the quotation marks rescue the statement from a potential problem (its falsity) and thus it would be silly to say that the statement is false meaning that these are not proofs. If pragmatic intrusion serves to rescue statement (13), then it procures immunity against a use of ‘That is false’ as a claim that these are not proofs. However, see what happens if truth-conditional pragmatic intrusion occurs and (13) is interpreted as (14):

(14) These are what John would call (or normally calls) ‘proofs’ of the illustrations.

If such a pragmatic intrusion is accepted, then it would be possible to reply:

(15) That’s false. John never used the word ‘proofs’.

Of course, this may well require a bit of stage setting. Suppose there are endless discussions in the office concerning John’s using the word ‘proofs’ in the past and that one or two employees often say things such ‘These are the proofs’ or ‘Send me the proofs’ and by using these expressions they allude to the fact that John used this expression. John, however, knows that it was not him, but Bill, that started using this expression in the office, but, for whatever reason, his colleagues started to circulate the information that it was John who did so. After noting that by the utterance ‘These are the proofs of the illustrations’, his colleagues actually allude to the information that John (and not Bill) used the term ‘proofs’, he may very well say: ‘It’s false. I never used the word ‘proofs’’. And by saying this, he would presumably reply to the proposition ‘This is what John would call ‘proofs’.’ Suppose now John is not sure whether his colleagues are actually alluding to him or, rather to Bill. Then he could say, Look if you are saying that this is what I would call ‘proofs’ you are wrong. This is what Bill would call ‘proofs’.

Consider now, another example by Predelli, like (16):

(16) Life is ‘what happens while you are making other plans’.

Suppose that it is assumed by the speaker and the hearer that it was John Lennon who uttered ‘what happens while you are making other plans’. Then a pragmatically enriched utterance of (16) could be

(17) ‘Life is what Lennon would describe as ‘what happens while you are making other plans’.’.

Now, if (17) is so blatantly a quotation, I believe that it is not unnatural to correct this statement saying: ‘That’s not correct, nobody really said this before’. Furthermore, one could take issue with some of the words used. The hearer may well suggest a replacement for ‘you’ such as ‘people’ and if (17) is recognized as a quotation, then the hearer may very well say something like ‘This is false; Lennon used the word ‘people’ not the word ‘you’.

At this point the discussion has become very theoretical and very much depends on whether or not we are inclined to accept that these quotations have a truth-conditional import such as (17).

However, we do not only have examples such as the ones above to account for. Those examples are probably not the most central ones. We also have examples such as:

(18) ‘House’ has five letters;

(19) He said ‘Go to hell’.

(20)) He said that I ‘almost killed her’

to account for. In examples such as (18) and (19), (20) falsehood is obtained as a result of replacing the words quoted with coreferential expressions. In fact, in a previous section on the applicability of the notion of explicature, I have already tested the idea that pragmatic intrusion is relevant to cases of quotation. It is amazing that the central cases are NOT discussed by Predelli, while he is more keen on discussing cases of free quotation where the locutionary verb is missing. However, even in these cases, if pragmatics is seen as reconstructing the missing locutionary verb, then quotation has a bearing on truth-conditional import.

Conclusion

In this chapter, contrary to the received view of quotation, I have gone for a thoroughly pragmatic process and I have dispensed with Recanati's view that quotation is partially due to conventional implicature, and partially due to conversational implicature (especially open quotations seem to be cases where pragmatic interpretations abound). I am persuaded that a radical pragmatic view is more parsimonious.

I am persuaded that quotation, in virtue of its pragmatic mechanisms, can be compared to a language game. It is a practice – a social practice – which very closely resembles the phenomenon which Goffman called 'framing'. Framing activities very often operate on the basis of stylistic information. Consider for instance a film, with many flash-backs. It is indispensable that we should be able to distinguish a series of events which unfold in chronological order from another. Marking the differences in terms of style is what allows viewers to recognize the framing activities. In the same way, a university lecturer may insert some windows in his lecture, connecting the lecture to the present time and the present occasion (the location, the type of audience) and of course, she faces the same problem that one faces in quotation: having to mark stylistically a temporary break of frame (the lecturer may mark such breaks of frame by assuming a more amusing tone of voice). Similarly, talk at work is marked by stylistic signals, such as, for example, using a dialect as opposed to using the standard language, which is reserved for the most institutional parts of the proceedings. Many examples of this type can be offered to readers to characterize the connection with quotation. Quotation is a *social practice*, which means that you need to be embedded in societal activities to see how the practice works and how to practice the practice. Very often, it is the specific language game one is involved in that serves to frame the quotational part of one's discourse. Definitions – which may occur at an early stage in a book, in a lecture or in an academic chapter – are part of a societal practice, in which speakers and the audience take their time to reflect on and define the terms that are being used more frequently and systematically. In the context of this *definitional language game*, it is clear that words lose their extensional value and that we need to mark certain portions of the utterance as referring to words or expressions. Quoting other people's words *verbatim* is also an important practice, which usually also involves reporting the context in which the

words occurred, as, deprived of their contexts, words are like orphans and are deprived of most of their force, which is acquired in the context of specific language games. Consider, for example, what happens if a certain headmistress, in the course of a meeting says: suppose I say 'teacher Buccheri is an idiot'. This fragment of reported speech cannot be reported in isolation from its context, because otherwise one does not know the kind of embeddings which the words undertook. A first framing activity is the one of supposition; a second framing activity is one of exemplification. Suppositions are speech acts which involve a marked departure from normal assumptions; exemplification is also an activity requiring a different frame (you make an example while assuming that the event exemplified is not real). Granting that the example, despite the various types of embeddings, does not lose much of its aggressive force and that a real complaint can be reasonably issued, nevertheless, in the course of a legal controversy, if one deletes the context in which the example was made, one could very well think that the headmistress was crazy. The practice of reporting must at least include the competence required in reporting the parts of context which could reveal the depth of the embedding or, so to say, the various levels of framing.

Can the notion that quotation is a *language game* lead us away from the idea that pragmatics is involved in understanding what game one is faced with? Not necessarily. On the one hand, one must learn the practice of citing context as well as words, when one reports an utterance. On the other hand, citing context can lead to various meaning augmentations. Since the point of the quotation must be inferred, it is clear that the point of the quotation is variable and is tied to what the reporters have in mind at the moment of the utterance. A pragmatic reconstruction is certainly required. Facts about the stylistic nature of the quotational language game, furthermore, need not be learned piecemeal, but may well be integrated in a general pragmatic competence which involves resorting to the pragmatic principle of Relevance to infer an M-implicature.