Intention, common ground and the availability of semantic content: a relevance-theoretic perspective

1. Decoding and inference in linguistic communication

By its very definition, communication involves the transmission of information from one agent to another; if the communicator had no means of making some piece of information available to an audience, there would *a priori* be no possibility for communication to ever take place. Admittedly, in most its instances in nature, communication is achieved through the use of a common underlying code which allows the straightforward transmission of a message to an audience. Take for example the classic case of honey bees: A bee lets its beehive know the whereabouts of a good supply of nectar by moving its body in some particular manner. In this setting, communication occurs through the encoding of the information that is to be communicated (i.e. the location of the nectar) into a specific 'dance', which the audience bees are able to decode through the implementation of an identical copy of the code that the communicating bee is using. From a first look, human verbal communication can also be treated with a similar rationale, according to which, the existence of a common linguistic code between two interlocutors suffices on its own to guarantee their successful communication. In such an account, all that the speaker has to do is encode her message into a natural-language sentence and utter it. The hearer will in turn decode the utterance's meaning and faithfully reconstruct in this way the speaker's original message.

Though it may seem intuitively adequate, this scenario actually falls short of holistically explaining verbal communication. Largely due to Grice's pioneering work (1989), it has been generally acknowledged that human communication is not entirely amenable to the aforementioned code-based treatment, but is essentially achieved through the expression and recognition of intentions. From this perspective, upon recognising the speaker's intention to communicate a message to him and on top of decoding her communicated utterance, the hearer needs to inferentially construct the speaker-intended meaning on the basis of the recovered decoded content and the context of utterance. In the Gricean tradition, the paradigm output of this inferential process is an *implicature*, a proposition that is purposefully conveyed by the speaker's utterance without being part of its decoded meaning and which the hearer is expected to infer for the speaker's communicative intention to be fulfilled.

In its original exposition, Grice's argument did not straightforwardly compromise the validity of the aforementioned code-based account of verbal communication, but rather added an inferential layer to the decoding process that was customarily thought to take place during the comprehension of an utterance. This essentially brings us to the currently most common understanding of the semantics/pragmatics distinction and, according to which, semantics is the study of the linguistic code alone and, thus, responsible for providing us with the public meaning of sentences while pragmatics deals with the inferential processes that are needed for the construction of the speaker-intended implicated meaning when sentences are uttered in context. It is with respect to this view that one of the most heated debates in modern linguistics has arisen. For quite some time now, several formidable minds have been engaged with questions regarding the necessity and magnitude of contextual intrusions in the delineation of an utterance's explicitly expressed propositional content and the debate has been perpetuated through the argumentation of scholars belonging to two conflicting camps of thought, which Recanati (2004) neatly locates in the traditions of *Literalism* and *Contextualism* respectively. Roughly speaking, the literalist contends that natural language sentences carry truth-conditional content on their own while the contextualist argues that it is only in the context of its utterance that a sentence can be assigned a determinate truth value. Effectively then, the point of disagreement between the two doctrines amounts to whether fully propositional semantic content can exist in isolation from contextualisation.

Naturally, in line with the general fascination of the contemporary linguist with formal approaches to semantics, it comes as no surprise that literalism is currently the dominant position with respect to the debate at hand. In their recent attack against contextualism, Cappelen and Lepore (2005) take the literalist position, which is clearly represented in their positive theory of Semantic Minimalism, to be only commonsense, given the classic Fregean claim that in communication we share thoughts. As they characteristically note, "if communicated contents are restricted to (or, essentially tied to) specific contexts of utterance, then it is hard to envision how speakers who find themselves in different contexts can communicate" (2005:153). Provided then that contexts are personal and cannot be shared in the code-like manner that semantic contents are hypothesised to be shared across interlocutors, Cappelen and Lepore view the proposition expressed by a sentence irrespective of its context of utterance as our sole "minimal defense against confusion, misunderstanding, mistakes" during verbal communication (2005:185). Therefore, Semantic Minimalism advocates that, in their communicative practices, both the speaker and the hearer mentally

access not only identical, but also fully propositional semantic contents of the utterances utilised.

One of the frameworks that Cappelen and Lepore consider to be in direct opposition to their regime is Relevance Theory (Sperber and Wilson 1995). This is of course hardly unjustified, since relevance theorists hold a radical version of the linguistic underdeterminancy thesis, according to which, "linguistically encoded meaning *never* fully determines the intended proposition expressed" (Carston 2002:49, emphasis in original). From the relevance-theoretic perspective, during comprehension, an utterance's semantic content might be indeed decoded in the aforementioned sense, but the outcome of this decoding process, the utterance's *logical form*, always needs considerable contextual input to gain full propositional status. In this respect, the inferential computation of the speaker's intentions while a logical form is recovered becomes crucial for the construction of an utterance's explicitly expressed content by the hearer and the product of this *parallel* and *simultaneous* operation of decoding and inference is a pragmatically enriched, yet properly truth-evaluable, proposition, which relevance theorists dub *explicature*.

Clearly, this approach is largely incompatible with Cappelen and Lepore's account. Its basic shortcoming, as the proponents of Semantic Minimalism themselves isolated it recently (Cappelen and Lepore 2007), is that it jeopardises a central aspect of verbal communication, that is, the relation of identity that, according to Cappelen and Lepore, needs to exist between the basic proposition expressed by the speaker and the one constructed by the hearer for communication between them to take place. Indeed, by allowing context to intrude in the delineation of an utterance's explicit content, "we need recognize only speaker-relative content and listener-relative content and a relation of *similarity* holding between these two contents" (Bezuidenhout 1997:198, emphasis in original). And this is a view that relevance theorists certainly share with most contextualists. For Cappelen and Lepore, however, this position is seriously flawed:

> […] there's no non-trivial sense of 'similarity' in which the explicatures arrived at by using the [relevance-theoretic comprehension procedure…] will be similar to the proposition that the speaker intended to communicate. They will, if R[elevance] T[heory] is correct, be developments of the same logical form LF, but an LF can be developed into *radically different* propositions. (Cappelen and Lepore 2007:117, emphasis in original)

Noticeably, this is hardly a new criticism. As Fodor and Lepore have also argued, "the kind of explanations that semantic theories are supposed to give would not survive substituting a

similarity-based notion of content for an identity-based notion" (Fodor and Lepore 1999:382), which is on its own particularly problematic since "nobody has the slightest idea how to construct the required notion of content similarity" (ibid).

Against this background, in the following section, I will defend the contextualist orientation by providing an overview of some basic sources of linguistic underdeterminancy, as these are pinpointed in the relevance-theoretic literature. Then, I will briefly present the basic tenets of Relevance Theory with a view to show that the comprehension procedure that it puts forth realistically addresses the on-line process by which interpretation takes place, safeguarding the validity of its outcomes in a non-trivial way. In the remainder of this paper, I will turn to the relevance-theoretic framework itself, evaluating further its conception of semantics. Discussing the implications that this conception carries with respect to lexical/conceptual content, I will finish my argumentation by proposing that relevance theorists need to recognise the context-sensitivity of semantic content even more cordially than they currently do.

2. The linguistic underdeterminancy thesis

The basic motivation of the contextualist stance in general and of the relevance-theoretic framework in particular is the need to provide a psychologically plausible account of the processes that are involved in the communication of utterances. Conversely, the roots of literalism are located in the early proposals of *ideal language philosophers*, like Frege, Russell and Tarski, who sought to investigate meaning from the perspective of formal logic and "were not originally concerned with natural language, which they thought defective in various ways" (Recanati 2004:1). For instance, Frege argued for "a sharp separation of the psychological from the logical, the subjective from the objective" (Frege 1997:90) which indeed rendered natural language defective for his purposes, language being clearly a "mixture of the logical and the psychological" (Frege 1997:243).

In this setting, formal approaches to semantics aim at examining necessary truths and essentially maintain that "being true is quite different from being held as true, whether by one, or by many, or by all, and is in no way to be reduced to it" (Frege 1997:202). In contrast, contextualist accounts are specifically concerned with the "intuitive truth-conditions" of an utterance (Recanati 2004:24), that is, with the conditions under which what an utterance asserts can be held to be true by its hearer. From this perspective, the context-independent truth-conditional content that formal semanticists assign to a sentence appears to be the result

of a "principle which has absolutely no bearing on human psychology" (Carston 1988:165). However, Cappelen and Lepore (2005:176-189) attempt to rebut this criticism by arguing that their minimal propositions are psychologically real and, therefore, considered by the hearer during the interpretation process. Reinstating the contextualist criticism, in what follows, I will illustrate how the explicitly expressed propositional content of an utterance is customarily much more context-dependent than Cappelen and Lepore wish to persuade us.

To begin with, even Cappelen and Lepore acknowledge that a proper theory of semantics needs to incorporate at least some minimal context-sensitivity. It is by now commonly accepted that there exists a number of linguistic expressions, customarily called *indexicals*, which are clearly context-dependent. Even intuitively, expressions like 'she', 'here', 'now' have a different truth-conditional contribution on most occasions of their use. Against this background, the basic tenet of Semantic Minimalism becomes that "the semantic content of a sentence S is the proposition that all utterances of S express (*when we adjust for or keep stable the semantic values of the obvious context sensitive expressions in S*)" (Cappelen and Lepore 2005:2-3, emphasis my own). Based on Kaplan (1989), Cappelen and Lepore provide a list of these obviously context-dependent indexicals, assuming that these are merely exceptions to the rule of semantic content autonomy. In my view, this is hardly the case, given the frequency with which we employ such expressions – especially pronominals – in our everyday communicative practices, but, in all fairness, I will grant Cappelen and Lepore their window of doubt, since there is a sense in which the context-sensitivity of indexicals is special. Even in the contextualist literature, indexical expressions are clearly distinguished, since in their case it is the lexical item itself that mandatorily triggers the inferential process by which it will be assigned a determinate semantic value. And this process, often called *saturation*, is qualitatively different from the corresponding process of *modulation*, that is, the context-driven process by which the semantic content of an utterance is freely enriched in order to gain full propositional status.

Nevertheless, even if we accept Cappelen and Lepore's argument that the linguistically triggered context-sensitivity of a small set of expressions is indeed exceptional, the problems for Semantic Minimalism are not eradicated. In order to show this, I will now provide a – necessarily schematic – summary of some basic sources of linguistic underdeterminancy, as these are discussed at length within the relevance-theoretic framework by Carston (2002:15-93).

In this respect, Semantic Minimalism may be prepared to accommodate the obvious

context-dependence of indexicals, but there are still further cases of linguistic expressions whose reference cannot be determined without resorting to the context of utterance and which Cappelen and Lepore do not include in their list of context sensitive expressions. For one, much like indexicals, proper names also refer to different individuals in different contexts (Sperber and Wilson 1995). Similarly, as Higginbotham notes (1988), the assignment of reference to 'incomplete' definite descriptions – as in (1) – and 'specific indefinite' descriptions – as in (2) – can correspondingly be problematic for traditional truth-conditional accounts of semantics:

(1) The president (whoever s/he is) is in trouble.

(2) If a certain blonde calls, pass her through directly to me.

Accordingly, even definite descriptions with no indexical element whatsoever can be taken to be context-dependent. As Recanati discusses (1987, 1996), in every communicative act there is a *domain of discourse*, "with respect to which the speaker presents his or her utterance as true" (Recanati 1987:62), and on which the reference of a definite description always depends. In order to illustrate this argument, let's consider the case of Ann's father uttering (3) to Ann during her graduation ceremony in 2008:

(3) Mum will be pleased to see that the current prime minister of Britain attended
your graduation.

Imagine now that both Ann and her father know that Ann's mother believes that Blair is still the prime minister of Britain when he is not. In this case, Ann's father uses the definite description 'the current prime minister of Britain' to refer to Blair, and Ann is respectively expected to interpret it with respect to her mother's belief system. In this particular setting then, the domain of discourse in which Ann will decode her father's utterance is not the actual world, as literalist accounts of semantics would predict. Along these lines, it becomes possible that the relativity of reference to domains of discourse can actually apply to all definite descriptions, even the ones that seem "the least likely to yield to the general context-dependence thesis" (Carston 2002:38). Finally, the context sensitivity of an utterance's propositional content is also apparent in the case of lexical or syntactic ambiguity. For instance, in the tradition of semantics, the problem of lexical ambiguity is usually tackled with by positing what Pustejovsky (1995) calls a *sense enumerative lexicon*, which comprises different lexical entries for each sense that an ambiguous lexical item is considered to have. In this setting, the context of utterance again plays the essential role of disambiguating which of the two semantic entries of, say, the word 'bank', its every use points to.

Furthermore, the linguistic string employed in an utterance might still fall short of determining a full proposition "even after all necessary reference assignments and disambiguations have taken place" (Carston 2002:22). Consider the examples in (4) to (8) that follow (ibid):

(4) Paracetamol is better.   [than what?]

(5) It's the same.   [as what?]

(6) She's leaving.   [from where?]

(7) He is too young.   [for what?]

(8) It's raining.   [where?]

What is obvious here is that for the proposition expressed by these utterances to be assigned a determinate truth value, certain missing elements, often referred to as *unarticulated constituents*, need to be contextually supplied. And regardless of whether these constituents are treated as hidden indexicals (Stanley 2000) or as contextually derived elements of the proposition expressed (Recanati 2002), the context manages in both cases, albeit in different ways, to intrude in the truth-conditional content of an utterance.

Another source of linguistic underdeterminacy can be found in instances where a lexical item's scope is left unspecified from the semantics of the item itself. In this setting, Gross (2001) discusses the 'part' context sensitivity of adjectival predicates, as this is exemplified in (9):

(9) The book is black.

It should be straightforward that in this case the meaning of 'is black' is semantically underdetermined in relation to the property of the book that it refers to (the cover, the dominant part of the cover, the pages, etc). And admittedly, context-sensitivity of this type carries over to other expressions as well, such as verbs, nouns etc., as is evident in (10):

(10) John finished the book.

Again, the semantics of 'finish' cannot on its own provide us with a full propositional content for this utterance, as John might have finished reading the book, writing it, binding it, and so on and so forth.

Finally, a related source of linguistic underdeterminacy can be located in instances where a lexical item's literal meaning, what semantic theory takes to be its encoded content, needs to be pragmatically adjusted for the communicator's intended meaning to be constructed by the hearer. Consider the cases in (11) and (12):

(11) Mary has a temperature.

(12) The fridge is empty.

It would be difficult to come up with the propositions explicitly expressed by these utterances in certain – and most likely familiar – contexts, if we did not contextually enrich the semantic content of 'temperature' and 'empty'. In (11), Mary's 'temperature' can be easily attributed the narrower interpretation of 'a high temperature' rather than its literal one, which would apply to Mary as it does to all living organisms. Similarly, in (12), the fridge will probably not be interpreted as being totally empty, but rather insufficiently filled with the goods that are needed by a household on a daily basis. In this case, the encoded content of 'empty' needs to be broadened for its more precise intended interpretation to be yielded.

Although my overview here has been inevitably brief, I believe that it can suffice to show that the underspecification of semantic content does not evaporate once indexicality is recognised. In this respect, the literalist argument for the existence of context-independent basic propositions that are identically shared among interlocutors becomes particularly shaky, as contextual intrusions are, on multiple occasions, indispensable for the determination of an utterance's explicitly expressed proposition. Therefore, endorsing the linguistic underdeterminancy thesis, at least from the relevance-theoretic perspective, is more than a merely contingent matter: "underdeterminancy is an essential feature of the relation between linguistic expressions and the propositions (thoughts) they are used to express" (Carston 2002:29).

3. The relevance-theoretic comprehension procedure

As the previous section showed, contextual intrusions in the proposition explicitly expressed by an utterance are so frequent that they practically prevent the hearer from recovering a basic proposition that will be identical in content to the one that the speaker intended to communicate. On the face of this impossibility, it becomes obvious that the relation of similarity that contextualism holds to exist between the speaker-uttered and hearer-constructed contents essentially constitutes the only realistic alternative in a psychologically plausible discussion of our communicative practices. Indeed, for the relevance theorist, it is "neither paradoxical nor counterintuitive […] that communication can be successful without resulting in an exact duplication of thoughts in communicator and audience" (Sperber and Wilson 1995:193). However, as I noted above, Cappelen and Lepore reject this position, for it leaves open the possibility that the hearer will construct a radically different explicature from the one that the speaker might have intended to communicate. In

this section, I will argue that even though the inferential process by which logical forms are developed into full propositions can in principle lead to this result, if we take the relevance-theoretic proposals seriously, it becomes highly unlikely that this will often happen in practice.

To begin with, it seems necessary to note that Relevance Theory is essentially a theory of cognition and mental processing and that this is how it manages to shed light on how verbal communication realistically takes place. Using the term in a technical sense, Sperber and Wilson explore relevance as a psychological property of cognitive input to mental processing. In the current setting then, this input can be identified with a communicated utterance that needs to be processed in order to construct its intended meaning. From the relevance-theoretic perspective, the degree in which an utterance will be relevant to its hearer, and thus worth processing, depends on a balance of cognitive effects and processing effort:

<u>Relevance of an input to an individual</u>
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.

While the notion of processing effort should be pretty self-explanatory in cognitive terms, the one of positive cognitive effects could do with some minimal elaboration. As should be evident by now, according to the relevance-theoretic framework, in order for the hearer to construct the message that the speaker intended to communicate to him, regardless of whether this message is explicitly expressed or implicated, he will need to utilise his general inferential abilities. In this respect, relevance theorists define non-demonstrative inference, as spontaneously used by us in our communicative practices, as a process that "starts from a set of premises and results in a set of conclusions which follow logically from, or are at least warranted by, the premises" (Sperber and Wilson 1995:12-13). Effectively, this set of premises refers to the context, which comprises a number of assumptions that the hearer brings to the forefront of his attention, and against which he will infer the speaker-intended meaning during the comprehension of an utterance. The conclusions derived from the unification of the utterance with the context of interpretation then, will achieve positive cognitive effects if they bring about an improvement in the hearer's system of beliefs, by

altering his already existing assumptions in a non-trivial way. Therefore, what follows from the above definition is that, the more an utterance improves the hearer's belief system and the less effort it requires in its processing, the more relevant it will be to him.

Having defined relevance in this way, Sperber and Wilson go on to discuss the notion's significance in relation to the way in which all mental computations take place and propose that considerations of relevance ultimately orchestrate the operation of our cognitive system in important respects. This is spelled out in their first or cognitive principle of relevance:

Cognitive Principle of Relevance
Human cognition tends to be geared to the maximization of relevance.

Following research in the domain of evolutionary psychology, Sperber and Wilson argue that, like most biological systems, human cognition should be viewed as a system that has evolved through processes of natural selection. Now, there is undoubtedly a vast number of stimuli surrounding us at any given moment, stimuli that can trigger our mental processing, but since we could not possibly attend to them all at the same time, we need to select the ones that we will focus on. In this respect, it is only in our interest to offer priority to those inputs from which we will benefit the most. In evolutionary terms, this creates an immediate pressure for our cognitive mechanism, a pressure to which it can be thought to have responded over time by adapting into a system that gives priority to those inputs that will provide it with the largest gains. And actually this certainly seems to be a valid observation if we consider the fact that certain stimuli, like the sound of an explosion, automatically impinge on our attention. Similarly, the need for an all the more efficient operation would correspondingly create a pressure for the mind to minimise the effort that it will have to spend in the processing of such particularly useful stimuli. Here, a case in point would be communicative stimuli, such as linguistic utterances, which we are normally able to comprehend, performing complex inferences, within only a few milliseconds. Against this background, it seems reasonable to assume that the human cognitive system has a natural tendency to attend to the most beneficial stimuli and allocate its resources in such a way that this processing will come as effortlessly as possible. And Sperber and Wilson's notion of *maximal relevance* can be directly implemented in the technical exposition of this argument, as it straightforwardly captures the idea that the search for the greatest possible effects for the least possible effort

has been instrumental in our cognitive system's evolution.

Turning to communication now, Sperber and Wilson discuss communicative stimuli employing the same cognitive definition of relevance and positing in their argumentation a dedicated inferential processor, whose domain is purposefully communicated stimuli, and which automatically computes the full set of cognitive effects that a communicative stimulus produces in the mind of the addressee. Assuming then that human cognition is indeed geared towards maximal relevance, as the first principle of relevance predicts, and that communicative stimuli automatically pre-empt our attention, an interesting picture emerges. For one, it makes sense to suppose that an addressee can automatically expect a communicative stimulus to provide him with some adequate cognitive effect or else he would have the choice of not processing it. Then, given that the central goal of communication is to understand and be understood, it seems to be "in the communicator's […] interest both to do her best and to appear to be doing her best to achieve this […] goal" (Sperber and Wilson 1995:268) by planning her utterance accordingly, otherwise the addressee's relevance-oriented cognitive system would again not necessarily pay attention to her communicative stimulus, as it automatically does. In the relevance-theoretic framework, these two observations combined create a fundamental constraint for the way in which an addressee will *always* treat a communicative stimulus, a constraint that is captured in Sperber and Wilson's second or communicative principle of relevance:

Communicative Principle of Relevance
Every act of ostensive communication communicates a presumption of its own optimal relevance.

In this principle, what the *presumption of optimal relevance* amounts to is merely what was established above, that is, that an ostensively communicated stimulus will always be treated by an addressee as 'relevant enough to be worth his effort to process it' and as being 'the most relevant one compatible with the communicator's abilities and preferences'. Therefore, since once we recognise a particular stimulus as *ostensively*, that is, deliberately, communicated to us, it is impossible for us to not process it, we will always process it as optimally relevant.

In this setting, a fundamental difference between maximal relevance that applies to cognition and *optimal relevance* whose presumption guides the processing of communicative stimuli arises: the former signifies the best possible balance between effort and effect while

the latter's goal is the retrieval of adequate cognitive effects for no unjustifiable processing effort expenditure. Naturally, during communication, the cognitive effects that a stimulus will produce in the mind of the addressee will be deemed adequate once his derived interpretation can be accepted as the one that the communicator intended him to arrive at. In this respect, the second principle of relevance predicts that, since hearers cannot but process communicative stimuli as optimally relevant, the comprehension procedure will end when these satisfactory cognitive effects are produced. And regarding the on-line process by which these cognitive effects will be sought by the addressee, Relevance Theory predicts that it will automatically take place following a specific route:

> Relevance-theoretic comprehension procedure
> a. Follow a path of least effort in computing cognitive effects: Test interpretive hypotheses in order of accessibility.
> b. Stop when your expectations of relevance are satisfied (or abandoned).

Admittedly, this comprehension procedure follows directly from the two principles of relevance that I have outlined above. For one, given the necessity of contextual intrusions in the delineation of an utterance's basic proposition, as the linguistic underdeterminancy thesis clearly demonstrates, if human cognition were not geared towards maximal relevance, a hearer would have to arbitrarily test an infinity of interpretive hypotheses before reaching the most satisfactory one; a task that he clearly has neither the time nor the cognitive resources to perform for every single communicative act. Therefore, accepting that his cognitive system will have a natural tendency to chunk particular pieces of information together in a relevance-boosting manner, so that he can go on and test these hypotheses in the order in which they become accessible to him, is "not just a reasonable thrift, [...but] an *epistemically* sound strategy" (Sperber & Wilson 1996:532). Accordingly, as the presumption of optimal relevance creates in the addressee the expectation that the communicator will have planned her utterance so that its meaning will be constructed by him without any unnecessary processing effort expenditure, it makes sense to assume that he will stop processing its meaning once he comes across an interpretation which he finds satisfactory with respect to the communicator's intentions along the path of least effort.

Against this background, it becomes evident that the relevance-theoretic comprehension procedure is not merely a convenient assumption regarding the interpretation of some communicative stimulus, but follows from a series of independently motivated

arguments regarding the evolution and organisation of the human cognitive system and the rationality that underlies human communication. In turn, these arguments safeguard that, on most occasions, the outcome of the interpretation process, be it some explicature(s) or implicature(s), will not substantially depart from the communicator's original intentions, although it might well do, as customarily happens in cases of miscommunication.

4. The public availability of lexical meaning

Up to this point, I have demonstrated how the relevance-theoretic approach can successfully react to the literalist tradition that takes semantic knowledge to be capable of providing an utterance's explicitly expressed propositional content on its own. Having addressed the procedure by which it foresees that the inferential developments of an utterance's logical form will lead the hearer to the speaker-intended meaning in a psychologically realistic manner, I will now turn to assess the way in which the relevance-theoretic framework identifies lexical semantic content in its premises.

In this respect, Sperber and Wilson seem to hold a rather traditional view of semantics, according to which, an utterance's decoded content, its logical form, can be identified with its underlying sentence's semantic representation:

> By definition, the semantic representation of a sentence, as assigned to it by a generative grammar, can take no account of such non-linguistic properties as, for example, the time and place of utterance, the identity of the speaker, the speaker's intentions, and so on. The semantic representation of a sentence deals with a sort of common core of meaning shared by every utterance of it. (Sperber and Wilson 1995:9)

From this statement, it becomes clear that the relevance-theoretic notion of logical form is essentially context-insensitive, which further suggests that Relevance Theory and Semantic Minimalism are not as incompatible as they might seem to be at first sight. As Wedgwood (2007) first noted and Carston (in press) asserts, both traditions pinpoint some semantic content upon which pragmatic inference effectively operates. Naturally, the relevance-theoretic conception of this encoded content departs substantially from the corresponding minimalist one, since its contextual development which occurs alongside its decoding, is indispensable for it to gain full propositional status.

However, no matter how schematic or propositionally incomplete a sentence's logical form might be considered from this perspective, it still comprises some content, the semantic

13

content of the lexical items that constitute it, that cannot but be virtually context-independent. And following Sperber and Wilson's rationale, this encoded lexical meaning should be respectively viewed as the common core of meaning shared by every usage of the lexical item in abstraction from individual contexts of use. Therefore, even though relevance theorists argue that the hearer "entertains thoughts [and not…] semantic representations of sentences" (Sperber and Wilson 1995:193) at the propositional level, they still maintain that at the lexical one he will need to entertain an expression's semantically encoded content before deciding whether to contextually enrich it or not. In the setting of psychological plausibility in which Relevance Theory purports to investigate linguistic communication then, a fundamental question immediately presents itself: How are we to realistically pinpoint what this context-independent meaning is with respect to some linguistic expression?

To begin with, it seems necessary to briefly illustrate the Fodorian background on which the relevance-theoretic conception of semantics squarely lies. In this respect, much like Fodor, relevance theorists distinguish between 'linguistic' and 'real' semantics in the following way:

> linguistic semantics […] could be described in statements of the form *'abc' means (= encodes) 'ijk'*, where 'abc' is a public-language form and 'ijk' is a Mentalese form (most likely an incomplete, schematic Mentalese form) [while…] 'real' semantics […] explicates the relation between our mental representations and that which they represent […] and whose statements may take the form *'hijk' means (= is true iff) such-and-such.* (Carston 2002:58)

In this sense, Relevance Theory adopts the further Fodorian argument that language "inherits its semantics from the contents of beliefs, desires, intentions, and so forth that it's used to express" (Fodor 1998:9). By definition then, the semantic content of a lexical item is inherited by the 'real' semantics of its associated mental concept. And regarding this 'real' semantics, Fodor adopts a purely externalist perspective, according to which, the true uses of a concept, such as CAT, are caused by actual cats.

Similarly, when addressing lexical meaning, Sperber and Wilson begin their discussion by positing mental concepts as the direct cognitive counterparts of lexical items. In this respect, "the 'meaning' of a word is provided by [its…] associated concept (or, in the case of an ambiguous word, concepts)" (Sperber and Wilson 1995:90), where a concept can be viewed as a stable address in memory comprising three entries. Of these entries, the most

straightforward to approach seems to be the *lexical* one, which is taken to encompass syntactic and phonological information about the lexical item that encodes a concept in natural language. Then, a concept's *logical* entry consists of a set of inferential rules, commonly known as *meaning postulates*, whose function is to provide the inferential processor with sets of premises and conclusions that "capture certain analytic implications of the concept" at hand (Carston 2002:321). Finally, the *encyclopaedic* entry of a concept contains general information that we individually hold with respect to its denotation; information that can be arbitrarily stored in the form of full propositions, assumption schemas or mental images.

From the relevance-theoretic perspective, the existence of both a lexical and a logical entry in the same conceptual address is deemed necessary given each concept's double contribution to decoding and inference during utterance interpretation. As Sperber and Wilson themselves observe, "recovery of the content of an utterance involves the ability to identify the individual words it contains, to recover the associated concepts, and to apply the [inferential…] rules attached to their logical entries" (Sperber and Wilson 1995:90). In this respect, it is the lexical entry of a word that activates the meaning postulates that are included in the logical entry of its encoded concept, which in turn enter the inferential processor one-by-one during an utterance's decoding into its logical form. Regarding the distinction of the logical from the encyclopaedic entry now, Relevance Theory argues that it essentially corresponds to the traditional analytic/synthetic distinction. Along these lines, a concept's logical entry is perceived to comprise inferential rules that express analytic, and thus necessary, truths with respect to the concept's meaning while its encyclopaedic entry contains chunks of memorised information that can optionally enter the inferential processor contributing to contextual enrichments of the concept's encoded content.

In this respect, it seems to me that the most salient way in which we can construe a lexical item's encoded content from the relevance-theoretic perspective is by identifying it with the inference rules that are included in its associated concept's logical entry. And the following argument, put forth by Sperber and Wilson themselves, certainly points to that very direction:

> Encyclopaedic entries typically vary across speakers and times: we do not all have the same assumptions about the Napoleon or about cats. They are open-ended: new information is being added to them all the time. There is no point at which an encyclopaedic entry can be said to be complete, and no essential minimum without which

one would not say that its associated concept had been mastered at all. Logical entries, by contrast, are small, finite, and relatively constant across speakers and times. There is a point at which a logical entry for a concept is complete, and before which one would not say that the concept had been mastered at all. (Sperber and Wilson 1995:88)

Therefore, by maintaining that a lexical item's semantic content is context-insensitive and thus stable across individuals and times, relevance theorists cannot but equate it with its associated concept's set of meaning postulates, since the information included in its encyclopaedic entry is highly individualistic and essentially private. However, given the aim of psychological plausibility with which Relevance Theory was developed to begin with, this seems like a rather counterintuitive move.

Following Quine (1951), it is now commonly accepted that purely logical analytic truths do not exist and that our intuitions of analyticity are empirically driven. In this respect, a statement might indeed appear to be *a priori* true, but is always amenable to revision, once further beliefs to which it closely relates are correspondingly modified. In light of this argument, even Fodor himself, who originally introduced meaning postulates in the discussion of conceptual content, was led to question their theoretical usefulness, eventually abandoning them from his theoretical investigation of concepts (Fodor 1998). Therefore, it seems that the conceptual content that relevance theorists identify as semantic crucially rests on an argument regarding analyticity that was ill-founded and psychologically unrealistic to begin with. Even so, Relevance Theory cannot discard meaning postulates as straightforwardly as Fodor did, since a central aspect of their account of utterance interpretation crucially rests on their very application. Without logical inferential rules, an utterance's decoding into its logical form would be practically unfeasible.

Against this background, the only way in which the relevance-theoretic framework could preserve its current account without compromising its fundamental aim of psychological plausibility is by substituting its traditional conception of meaning postulates with a more realistic counterpart. Such an attempt has been recently made by Horsey (2006), who proposes a replacement of the problematic notion of analyticity with that of *psychosemantic* analyticity, according to which, "while the majority of our psychologically represented inference rules are no doubt veridical, this is by no means necessary" (Horsey 2006:74). And given Quine's argumentation, it certainly becomes clear that the relevance-theoretic framework needs to incorporate psychosemantic analyticity so construed in its premises.

However, by necessarily psychologising analyticity, Relevance Theory cannot but respectively psychologise encoded conceptual contents, abandoning its view of lexical semantics as effectively context-independent. In the original relevance-theoretic account, it was only by means of the externalist scenario with respect to the way in which we acquire content-constitutive inferential rules that the identity of semantically encoded contents across individuals and times could be guaranteed. In this setting, if the logical entry of the concept CAT would include the inference rule 'ϕ CAT ψ → ϕ ANIMAL ψ', it would follow that the context-independent semantic content of the word 'cat' would nomologically contain at least the information that cats are animals in all its mental instantiations across individuals. By accepting psychosemantic analyticity, however, there can be plenty of cases where two individuals' meaning postulates with respect to the same concept will vary. To use an example that Horsey (2006:74) himself discusses, in the new picture, we can accept the meaning postulate 'ϕ WHALE ψ → ϕ FISH ψ' as content-constitutive of Mary's concept WHALE, even though it is not veridical in relation to the real world. However, there will certainly be other individuals whose logical entry of the same concept will not contain the very same meaning postulate, substituting it for 'ϕ WHALE ψ → ϕ MAMMAL ψ'. In this respect, it becomes evident that, if a concept's meaning postulates do not express necessary truths about the concept's extension in the actual world, but rather constitute psychological constructs that capture what an individual perceives to be necessary truths regarding it, then by definition the logical entry of this concept cannot be publicly shared. Therefore, since lexical semantic contents can typically vary across individuals, the very notion of some 'common core of meaning shared by every usage of a lexical item in abstraction from individual contexts of use' that Relevance Theory endorses becomes practically vacuous. Naturally, this does not necessarily entail that there is no lexical semantic content as such, but rather that, much like its propositional counterpart, this content also succumbs to the radical linguistic underdeterminancy thesis in the communicative setting.

In order to work out the main implication that this conclusion carries for the relevance-theoretic account, I have to first briefly illustrate the on-line process by which Relevance Theory addresses contextual enrichments at the lexical level. Following Barsalou (1983, 1987) and his experimental evidence regarding the flexibility with which we can entertain temporary mental constructs that arise in particular contexts, relevance theorists suggest that during utterance interpretation an encoded concept might get pragmatically adjusted forcing us to construct an *ad hoc* concept in its place. In this respect, after the

decoding process provides the inferential processor with the meaning postulates of the concept associated with a word, the inferential processor might either accept this encoded meaning as it stands or contextually adjust it, following the relevance-based comprehension procedure, in an attempt to provide certain positive cognitive effects that will satisfy the hearer's expectation of optimal relevance.

In this sense, from the relevance-theoretic perspective, all the lexical meanings that the hearer will construct during the interpretation of an utterance can be viewed as "outcomes of a single pragmatic process which fine-tunes the interpretation of virtually every word" (Wilson 2004:344). Given then the current proposal that the semantic content of a linguistic expression is best viewed as individualistic and cannot be publicly shared, the tantalising possibility that all communicated lexical meanings essentially correspond to *ad hoc* concepts seems to present itself. And that is because the hearer cannot but contextually enrich a lexical item's semantic content against the speaker-intended context, as if he does not do so, there is always a chance that his individualistic semantic content for that lexical item will not be similar enough to the one that the speaker intended to communicate. Even though this is a position that like most relevance theorists I am currently reluctant to explicitly defend, since it would require a complete reidentification of the theoretical notion of logical form, it certainly seems to provide fruitful ground for further investigation (for a more detailed argument along these lines, see Assimakopoulos 2008). For the time being, I believe that my argumentation should suffice for my current purposes of showing that lexical semantic content is much more context-sensitive than Relevance Theory currently maintains. Against this background, it seems that relevance theorists should eventually open up to the possibility of incorporating a more radical version of contextualist semantics in their framework's basic exposition.

5. Concluding remarks

One of the basic assumptions of most accounts of communication is that for communication to be successful some sharedness of information between the speaker and the hearer is necessary. Likewise, from the relevance-theoretic perspective, "communication requires some degree of co-ordination between communicator and audience on the choice of a code and a context" (Sperber & Wilson 1995:43). A basic innovation of the relevance-theoretic approach with respect to this issue has been that it does not take this co-ordination to be a given as traditional code theories of communication do. Rather, it takes the most

plausible route that interlocutors inferentially establish their common ground during their communicative practices. As I have shown in this paper, inference does not realistically kick in after the identification of some minimal semantic content that is publicly available to each and every one of us, but essentially operates *in parallel* to this decoding process, in ways that might be even more radical than relevance theorists currently acknowledge. Indeed, by endorsing the literalist perspective and its corresponding idea of semantic content identity between interlocutors, we could in principle explain communication rather easily. However, in doing so we would essentially turn our back on psychological plausibility; prioritising psychological explanation over theoretical convenience, it becomes crucial that we start to seriously entertain the possibility that discussions regarding the inferential attribution of intentions might be much more instrumental for the study of linguistic semantics than customarily perceived.

References:

Assimakopoulos, Stavros. 2008. Logical Structure and Relevance. PhD diss. University of Edinburgh.

Barsalou, Lawrence W. 1983. Ad hoc categories. In *Memory and Cognition* 11, pp. 211-227.

Barsalou, Lawrence W. 1987. The instability of graded structure. Implications for the nature of concepts. In Neisser, Ulric (ed.) *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*. Cambridge: Cambridge University Press, pp. 101-140.

Bezuidenhout, Anne. 1997. The communication of *de re* thoughts. In *Noûs* 31, pp. 197-225.

Cappelen, Herman and Ernie Lepore. 2005. *Insensitive Semantics. A Defense of Semantic Minimalism and Speech Act Pluralism*. Malden: Blackwell.

Cappelen, Herman and Ernie Lepore. 2007. Relevance theory and shared content. In Burton-Roberts, Noel (ed.) *Pragmatics*. Basingstoke: Palgrave Macmillan, pp. 115-135.

Carston, Robyn. 1988. Implicature, explicature, and truth-theoretic semantics. In Kempson, Ruth M. (ed.) *Mental Representations. The Interface between Language and Reality*. Cambridge: Cambridge University Press, pp. 155-181.

Carston, Robyn. 2002. *Thoughts and Utterances. The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Carston, Robyn. In press. Linguistic communication and the semantics/pragmatics distinction. To appear in *Synthese*.

Fodor, Jerry A. 1998. *Concepts. Where Cognitive Science Went Wrong*. New York: Oxford University Press.

Fodor, Jerry A. and Ernie Lepore. 1999. All at sea in semantic space. Churchland on meaning similarity. In *The Journal of Philosophy* 96, pp. 381-403.

Frege, Gottlob. 1997. *The Frege Reader (edited by Michael Beany)*. Oxford: Blackwell

Grice, Paul. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Gross, Steven. 2001. *Essays on Linguistic Context-sensitivity and its Philosophical Significance*. New York: Routledge.

Higginbotham, James. 1988. Contexts, models, and meanings. A note on the data of semantics. In Kempson, Ruth M. (ed.) *Mental Representations. The Interface between Language and Reality*. Cambridge: Cambridge University Press, pp. 29-48.

Horsey, Richard. 2006. *The Content and Acquisition of Lexical Concepts*. Unpublished PhD thesis, University College London.

Kaplan, David. 1989. Demonstratives. In Almog, Joseph, John Perry and Howard Wettstein (eds.) *Themes from Kaplan*. New York: Oxford University Press, pp. 481-563.

Pustejovsky, James. 1995. *The Generative Lexicon*. Cambridge, MA: MIT Press.

Quine, Willard V. 1951. Two dogmas of empiricism. In *The Philosophical Review* 60, pp. 20-43.

Recanati, François. 1987. Contextual dependence and definite descriptions. In *Proceedings of the Aristotelian Society* 87, pp. 57-73.

Recanati, François. 1996. Domains of discourse. In *Linguistics and Philosophy* 19, pp. 445-475.

Recanati, François. 2002. Unarticulated constituents. In *Linguistics and Philosophy* 25, pp. 299-345.

Recanati, François. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.

Sperber, Dan and Deirdre Wilson. 1995. *Relevance. Communication and Cognition (2nd edition)*. Oxford: Blackwell.

Sperber, Dan and Deirdre Wilson. 1996. Fodor's frame problem and relevance theory. In *Behavioral and Brain Sciences* 19, pp. 530-532.

Stanley, Jason. 2000. Context and logical form. In *Linguistics and Philosophy* 23, pp. 391-434.

Wedgwood, Daniel. 2007. Shared assumptions. Semantic minimalism and relevance theory. In *Journal of Linguistics* 43, pp. 647-681.

Wilson, Deirdre. 2004. Relevance and lexical pragmatics. In *UCL Working Papers in Linguistics* 16, pp. 343-360.