On Encoded Lexical Meaning: Philosophical and Psychological Perspectives *

Stavros Assimakopoulos †
stavros.assimakopoulos@um.edu.mt

ABSTRACT

The past few years have seen quite a bit of speculation over relevance theorists' commitment to Fodorian semantics as a means to account for the notion of encoded lexical meaning that they put forth in their framework. In this paper, I take on the issue, arguing that this view of lexical semantics compromises Relevance Theory's aim of psychological plausibility, since it effectively binds it with the 'literal first' hypothesis that has been deemed unrealistic from a psycholinguistic viewpoint. After discussing the incompatibility of Fodor's philosophical account with the perspective that relevance theorists adopt, I briefly suggest ways in which further behavioural research on the semantics/pragmatics distinction could help advance more cognitively-oriented accounts of encoded lexical meaning.

Keywords: lexical pragmatics, Fodorian semantics, relevance theory, "literal-first" hypothesis, polysemy.

* While preparing this paper I have benefited immensely from discussions with Lewis Bott, Ronnie Cann, Robyn Carston, Billy Clark, Felicity Deamer, Albert Gatt, Alison Hall, Anne Reboul, Nicola Spotorno, John Tomlinson Jr, Daniel Wedgwood and Deirdre Wilson. Of course none of the above is to be held responsible for the views expressed in this paper, which I would after all expect many of them to find objectionable. I am also particularly indebted to Francesca Ervas for her kind invitation to this special issue, editorial guidance, and, above all, incredible patience with me. The research presented herein has been financially supported by a travel grant provided by the ESF Euro-XPrag Network for collaborative research with John Tomlinson Jr, Lewis Bott and Julio Santiago de Torres, under the title "Varieties of lexical adjustment", as well as the Junta de Andalucía project entitled "Contemporary Naturalism and Pragmatism" (P08.HUM.04099).
† Institute of Linguistics, University of Malta.

Introduction

Ever since its inception in the seminal philosophical work of Paul Grice (1989), contemporary linguistic pragmatics has been typically viewed as an add-on to the more fundamental study of semantics, which deals with linguistic meaning in abstraction of its context of use. In this view, the recognition of speaker intentions, albeit instrumental for the recovery of 'what is implicated' by the utterance of a sentence, has been supposed to play little, if any role in the delineation of its semantic content, i.e. what Grice is customarily taken to have originally referred to as 'what is said' by it.

This way of distinguishing between semantic content and pragmatic import may still reflect the received view in the work of most contemporary researchers who are interested in linguistic meaning, but it has been repeatedly attacked in the past few decades. On a number of occasions, various scholars[1] have pursued the argument that, even over and above the obvious context-sensitivity of indexical expressions (such as *I*, *you*, *here*, *now*, etc), 'what is said' by a sentence cannot always be identified in isolation from the context of its utterance. Among them, relevance theorists have right from the beginning (at least as far back as Wilson & Sperber, 1981) insisted that the linguistically encoded meaning of an utterance, that is, its underlying sentence's semantics, falls short of determining the proposition explicitly expressed by it, and that the hearer has to undertake processes of pragmatic inference in order to work this proposition out. The thorough investigation of this claim primarily by Robyn Carston in a number of publications, culminating in her *Thoughts and Utterances* (2002), has eventually led Relevance Theory (henceforth RT) to take up a radical version of it, according to which, "linguistically encoded meaning *never* fully determines the intended proposition expressed" (Carston 2002, p. 49, emphasis in original).

The espousal of this position, which has been dubbed the *linguistic underdeterminacy thesis*, has obviously placed RT in direct opposition to the traditional way of carving the semantics/pragmatics distinction at the propositional level. Quite predictably, this has in turn sparked the reaction of various philosophers of language, who have criticized this deviation from the long-established way of studying linguistic meaning, counter-proposing

[1] See, for example, Searle (1978), Travis (1981, 1997), Sperber & Wilson (1986/1995), Carston (1988, 1999, 2002, 2004, 2008), Atlas (1989, 2005), Recanati (1989, 2001, 2002, 2004, 2010), Bach (1994a, 1994b, 1999), Levinson (2000), Jaszczolt (2005).

minimalist theories of semantic content that defend the context-independence of the proposition explicitly expressed by an utterance (e.g. Borg, 2004; Cappelen & Lepore, 2005). However, even minimalists themselves have at times recognized that "if communication can be shown to proceed without hearers processing the literal meaning of the sentence, i.e. without grasping minimal propositions, then the claim that minimal propositions have a unique role to play in actual communicative exchanges is undermined" (Borg, 2007, p. 353). In this respect, the linguistic underdeterminacy thesis seems to be vindicated by the existence of a quite substantial body of psycholinguistic research which has shown that the comprehension of figurative language does not necessarily involve the prior processing of the surface literal meaning (for representative overviews, see, Gibbs, 1994; Glucksberg, 2001).

Moving from propositional to lexical meaning, it is this very psychological implausibility of the so-called '*literal first*' hypothesis that has also motivated, as Deirdre Wilson herself notes (2011, p. 15), one of the latest developments within RT, that is, its account of ad hoc concept construction. Based on experimental research which has shown that our categorisation behaviour is highly context-dependent (e.g. Barsalou, 1987, 1992), relevance theorists have advanced the argument that the concept communicated through the use of some particular lexical item can be distinct from the concept it encodes, and thus requires a spontaneous process of pragmatic enrichment to be reached at during interpretation.

Even though the experimental investigation of the particular proposals that relevance theorists have put forth with respect to lexical pragmatics is still very limited, I will attempt in this short paper to show that their theoretical proposals have far-reaching implications for the discussion of word meaning, to the extent that they could even challenge the current view of encoded lexical meaning within RT itself. To this end, I will start off with a brief overview of the framework's assumptions that are relevant to the present discussion and will then move on to assess the account of lexical pragmatics put forth by relevance theorists and the tension it creates for traditional approaches to lexical semantics and, more specifically, the Fodorian one that relevance theorists have adopted right from the beginning. Wrapping up this paper, I will consider how the present argumentation can motivate new directions for behavioural research on the semantics/pragmatics interface.

1. Relevance-theoretic assumptions

Right from its emergence, RT has aimed at providing a thoroughgoing cognitive account of utterance interpretation. For this reason, relevance theorists draw the semantics/pragmatics distinction in terms of the different kinds of mental processing they take each of these types of meaning to be the output of, instead of discussing them at a theoretically abstract level, as most philosophers have so far tended to. From their perspective, semantic content is provided via decoding, which is performed by an autonomous linguistic mental module, while pragmatically derived meanings are taken to be generated by an inferential processor, which is in turn dedicated to the comprehension of deliberately communicated stimuli and effectively integrates the output of decoding with readily available contextual assumptions in the interest of calculating a reasonable hypothesis about the original speaker-intended meaning. Without getting into too much detail, which is after all unnecessary for my current purposes, RT predicts that a hearer will automatically comprehend a deliberately communicated utterance by following a path of least effort, according to which, he will assess interpretive hypotheses in order of accessibility until his expectations for an interpretation that will uncover the speaker's intended meaning are satisfied (or, in the case of miscommunication, abandoned).

In the current setting, the crucial aspect of the RT account is that it does not take inference to work on the overall output of decoding during the comprehension of a single utterance, as a traditional Gricean approach would have it; rather, the two modules work simultaneously, with the decoding one feeding input to the inferential every step of the way during the processing of the linguistically encoded stimulus. Obviously, the replacement of Fregean-style thoughts with subjective and context-dependent propositions – what relevance theorists call *explicatures* – makes it tempting to assume, as semantic minimalists have on occasion, that the framework encompasses some radically contextualist notion of semantics in its premises. On closer inspection, however, it turns out that this is not the case. As Daniel Wedgwood (2007) extensively discusses, the sole difference between the minimalist's way of describing semantic content and the relevance-theoretic notion of 'encoded meaning' lies in the contention of the former that sentences do actually encode full propositions; other than that, encoded meaning is equally 'properly'

semantic, in the traditional sense, for relevance theorists too. Considering the following passage from *Relevance*, this conclusion seems to be warranted:

> 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 & Wilson, 1986/1995, p. 9)

This identification of encoded meaning as essentially context-independent within RT also follows from the adherence of relevance theorists to Fodorian semantics. As Carston asserts (2002, p. 58), RT follows Fodor in assuming that 'real' semantics deals with the explication in truth-conditional terms of the relation between our mental representations and that which they represent, while linguistic semantics is merely 'translational', in the sense that public-language forms inherit their meaning directly from the Mentalese forms they encode. Therefore, since in the resulting picture the only actual bearer of semantic content is Mentalese, which consists in concepts, and, as Jerry Fodor has it, a concept is an unanalysable, monolithic atom that is individuated by some property of the real-world entity to which it is nomologically locked, it becomes virtually impossible for encoded semantic contents to vary across different contexts.

Therefore, even though relevance theorists maintain that sentences cannot be attributed any fully propositional semantics, they cannot but accept that lexical items[2] do encode context-independent meanings; that is, the ones that they directly inherit from their associated atomic concepts.[3] And this is indeed what they seem to have had in mind ever since they first entertained the idea that "words in a language can be used to convey not only the concepts they encode, but also indefinitely many other related concepts to which they might point in a given context" (Sperber & Wilson, 1998, p. 197), since, in the standard RT picture, in order for such ad hoc concepts to be constructed, the corresponding encoded concepts crucially need to be used as a starting point.

In order to briefly illustrate the rationale behind the RT account of ad hoc concept construction now, let's consider the following examples:

[2] Much like most relevant discussions, the present one deals exclusively with monomorphemic 'open-class' lexical items, i.e. words that have some descriptive content (unlike indexicals, connectives and the like).

[3] For a recently presented, yet still speculative, alternative view, see Carston (2012).

(1) John has a temperature.
(2) The fridge is empty.

It should be pretty straightforward that in order to come up with the proposition explicitly expressed by the utterances (1) and (2) in certain, most likely familiar, contexts, the hearer would have to somehow adjust the encoded meaning of 'temperature' and 'empty'. In (1), the word 'temperature' could be easily attributed the interpretation of 'a high temperature' rather than its actual denotational content provided by the concept TEMPERATURE, as this can be taken to be used in utterances like 'Celcius is a scale for temperature measurement'. Similarly, in (2) the fridge might not be interpreted as being totally empty, but rather as being insufficiently filled with the goods that are needed by a household on a daily basis. In this case, the encoded concept EMPTY would again need to be adjusted so that the 'not entirely empty, but insufficiently full' interpretation can be yielded.

According to the RT account, the construction of ad hoc concepts, like TEMPERATURE* and EMPTY* (to use their common notation), is the outcome of two pragmatic processes that can either individually or in unison contextually adjust the meaning that a lexical item's encoded concept carries. The first one, which is dubbed *narrowing*, results in meanings that are typically more specific than the encoded ones, such as the one of TEMPERATURE* in (1), while the second, *broadening*, respectively generates looser word interpretations. Apart from general approximation cases, like the one of EMPTY* exemplified above, concept broadening is also assumed to mediate the interpretation of hyperboles and metaphors, as well as category extensions (e.g. when a brand name, like 'typex' is used as an umbrella term for all products with a common function), neologisms and word coinages.[4]

As I have already noted in the previous section, the RT account of lexical pragmatics was originally motivated by Lawrence Barsalou's behavioural research on conceptual categorisation, a point that in itself gives the account a quite high degree of psychological plausibility. However, it has often been noted[5] that this plausibility is compromised when it is coupled with Fodor's philosophical semantics, which relevance theorists have adopted for the

[4] For detailed overviews of the RT account of lexical pragmatics, see Carston (2002: chapter 5), Wilson (2004), Wilson & Carston (2007).
[5] See, for example, Vicente (2005), Burton-Roberts (2007), Groefsema (2007), Assimakopoulos (2008), Reboul (2008), Vicente & Martinez Manrique (2010).

purposes of describing what gets decoded and fed into the inferential processor during interpretation. As Anne Reboul puts it, "the notion of an ad hoc concept has rendered visible a long-standing tension in Relevance Theory, viz. that between the adoption of an (atomistic and externalist) view of concepts (such as Fodor's) and the description that Relevance Theory effectively gives to concepts" (2008, p. 523). In the following section, I will turn to this tension and will attempt to show that the argument that relevance theorists have used against its existence runs into trouble when psychological considerations enter the picture.

2. Ad hoc concepts and the literal-first hypothesis

Based on descriptions similar to the one provided above, albeit much more detailed of course, Vicente (2005) and Groefsema (2007) have justifiably noted that the formation of communicated lexical meanings by means of broadening and narrowing suggests that encoded concepts must have some kind of internal structure, as it is only by way of manipulating such a structure that the construction of a speaker-intended ad hoc concept can be made possible. Naturally, if this is the case, encoded concepts cannot be conceived of as Fodorian atoms to begin with, since it is by definition impossible to either 'narrow down' or 'loosen up' a non-decomposable atom. This criticism seems to be further motivated by the description of concepts within RT itself, according to which, logical, lexical and encyclopaedic information is standardly assumed to be stored in different entries of a conceptual address. Given the explicit recognition of such different types of information associated with a concept and their implementation in various discussions of ad hoc concept construction, like, for example, when Carston suggests that, in narrowing, an encyclopaedic property of a lexically encoded concept can be 'elevated' to a logical (or content-constitutive) status (2002, p. 339), it certainly becomes all the more tempting to assume that encoded concepts must have some more substantive content than a monolithic atom would normally allow for in order for ad hoc concepts to be constructed on their basis.

The way in which RT addresses this criticism, however, can be easily located in Carston's parallelism (2010a, pp. 174–175, fn.6) of the relevance-theoretic notion of a conceptual address with what Fodor has recently called a 'mental file':

> When you are introduced to John [...] you assign him a Mentalese name and you open a mental file, and the same Mentalese expression (M(John)) *serves both as John's Mentalese name and as the name of the file that contains your information about John*; [...] according to this story, *we think in file names*; tokens of file names serve both as the constituents of our thoughts and as the Mentalese expressions that we use to refer to things we think about. (Fodor, 2008, pp. 94–95, emphasis in original)

Taking into consideration Fodor's description, it becomes clear that from the corresponding RT perspective, the various kinds of information that are thought to be associated with a concept do not form part of its semantic content per se, and thus play no role whatsoever in the decoding process as far as relevance theorists are concerned. Consequently, as a theoretical construct, an ad hoc concept would not appear to pose any particular problems for the way in which RT views semantics, as it is essentially the output of inference, with the input from decoding being solely the respective atomic concept (i.e. the Fodorian mental file name).

This line of argument, which Carston (2010a, 2010b) has followed in response to a slightly different, but comparable critique that Vicente & Martínez Manrique (2010) have put forth, is certainly reasonable when it comes to the deflation of the argument that relevance theorists would be better off employing a decompositional picture of lexical semantics rather than Fodor's atomistic account in their framework; yet, when the overall RT aim of developing a cognitively realistic account of communication is taken into consideration, it seems to be binding relevance theorists with a view that, as we have seen, they otherwise explicitly seek to distance themselves from, i.e. the 'literal-first' hypothesis.

Given the current RT account of ad hoc concept construction, according to which, the inferential enrichment of encoded concepts is standardly treated as an optional, top-down process, if decoding provides the inferential processor with the content of a word's encoded concept, this content cannot but be the first interpretive hypothesis that the hearer will test for relevance during the comprehension procedure. Consider, for example, the meaning communicated by the word 'temperature' during the interpretation of 'John has a temperature' in the aforementioned context in which John has a fever. Here, according to the current RT view, the output of the decoding of 'temperature' would be the concept TEMPERATURE, which carries the real, context-insensitive, and hence literal semantic content associated with the

word. For the inferential processor to construct the relevant ad hoc concept TEMPERATURE*, it will need to do so after testing the content of this encoded concept as a plausible hypothesis about the speaker-intended meaning, since, being by definition an unanalyzable atom, TEMPERATURE will necessarily be wholly employed in the process. But if the comprehension of figurative language is equivalent to that of literal meaning in processing terms, as the relevant experimental evidence suggests, the priority of this encoded meaning over the pragmatically enriched one is seriously compromised.[6] Given the mechanics of the cognitive systems that RT posits, even Wilson's recent suggestion that "the concept encoded by a word is *activated* during comprehension, but not necessarily *deployed* (2011, p. 16, emphasis in original) seems unsatisfactory, since again RT currently has no way of accounting for the activation of a concept without its initial incorporation in (and, if deemed unsatisfactory, potential discarding from) the mental representation that the inferential processor calculates as an utterance's basic explicature. In general, if the encoded denotational content of any concept can be bypassed during comprehension, it follows that the inferential processor has some way of discarding 'irrelevant' lexical meanings before actually assessing them as intended interpretations. But since the inferential processor's task is precisely to carry out this assessment in the first place, it has no way of knowing beforehand which encoded concept it will eventually keep intact and which it will need to enrich into an ad hoc concept.

If this line of reasoning is on the right track, it reveals a challenge that RT would need to tackle in order to satisfy its overarching aim of psychological plausibility. And while the relevant literature has focused almost exclusively on the implications that the account of ad hoc concept construction carries for the discussion of a concept's internal composition (or lack thereof), I believe that an equally important question that needs to be addressed is how well the philosophical discussion of lexical meaning that RT clings to can fit its

[6] A potential counter-argument that has been brought to my attention is that the experimental evidence against the literal-first hypothesis only carries implications for the discussion of propositional and not for that of lexical meaning. I think this is highly debatable since the metaphors used in the relevant literature often consist of a topic followed directly by the metaphor vehicle, as in the case of 'Her surgeon was a butcher' or 'My job is a jail'. According to the current RT account of lexical pragmatics then, it is only the concepts BUTCHER and JAIL that would need to get enriched for the figurative interpretation to become available, which effectively means that metaphor interpretation pertains more to the discussion of communicated meaning at the lexical rather than the propositional level (in at least such cases).

psychological orientation. As I will now turn to argue the two perspectives are quite hard to reconcile, and for good reason.

3. Philosophy and the psychology of encoded lexical meaning (all too briefly)

As we have already seen, much like Fodor, relevance theorists take on the commonplace assumption that semantic theory aims at providing an account of linguistic meaning at a level of abstraction from its actual use; 'a sort of common core of meaning shared by every utterance of a sentence'. Given the RT on-line processing picture, however, it is only natural to expect that the inferential processor will have already enriched bits and pieces of an utterance by the time it has been fully heard in an actual communicative setting. In this respect, even though the decoding of this utterance will generate a concatenation of context-independent conceptual representations, by the time an utterance's explicitly expressed meaning is constructed, it will inevitably present various degrees of deviation from the type proposition that Fodor's semantic theory puts forth. Following this rationale, Carston recently observed that even if a sentence's encoded meaning did somehow turn out to typically express a full proposition like the minimalist holds, the repercussions of this discovery would not be "a devastating blow for the central tenets of RT", as "the propositions concerned would usually be very weak/general or absurdly strong, often either truisms or obvious falsehoods", which would "almost never be the sort of contents that speakers want to communicate" (2010b, p. 268).

Indeed, considering the particular cognitive processing that mediates the comprehension of linguistic stimuli against the traditional philosophical context of studying semantics, Carston's remark appears to be on the right track, but from the very same psychological perspective, a pressing question also arises: if, without any contextual input, the thoughts that 'there are cats' or that 'it's raining' that Fodor alludes to in his discussions are never 'the sort of contents that speakers want to communicate', is there any principled reason for which we need to accept that they are thoughts that we ever even entertain? And if the answer to this question is negative, as I think it is, what is the reason for which we need to maintain that these semantic contents are actual thoughts - rather than artificial examples pertaining to an abstract model of thought - to begin with?

Turning to lexical semantic content, this problem becomes even more obvious. As Carston herself discusses (2002, p. 360),

> Focusing on the word 'happy', let's consider the concept that it is supposed to encode, a concept which is to provide communicative access to a wide range of other more specific concepts [...]. The idea is that the lexically encoded concept HAPPY is distinct from all of these; it is more general and abstract than any of them, but provides the basis, in appropriate contexts, for processes of pragmatic enrichment so that addressees can come to grasp one of the more specific concepts and incorporate it into their representation of the speaker's thought. But what is not at all clear is whether we ever actually have (hence sometimes try to communicate) thoughts in which this very general lexicalized concept features as a constituent, or indeed what the property of being HAPPY is, as opposed to being HAPPY* or HAPPY**, etc.

Clearly, this worry is not exclusive to 'happy', but rather seems to present itself when the encoded meaning of any gradable adjective, where no absolute denotational property exists, or even that of commonly used verbs like 'open' or 'stop' are put into scrutiny.[7] Even when we turn to nouns, the postulation by relevance theorists of ad hoc concepts in the mind, raises important questions regarding their implementation in the individual's everyday thinking too. So, when the doctor thinks that John has a temperature in our familiar by now context, is she thinking that he has a TEMPERATURE or rather a TEMPERATURE*? Similarly, when Mary thinks that she wants to meet a bachelor, to use another well-worn example from the RT literature, does she implement in her mental processing the concept BACHELOR, whose denotation includes all male individuals who are not married, or the narrower concept BACHELOR*, whose denotation includes those male individuals who are not married, but who would also be eligible candidates for her to marry (obviously not the Pope or some very old or gay man)? If, as I take it, the answer to these questions points to the ad hoc rather than the encoded concept, it follows that these concepts coexist in our conceptual repertoire alongside their Fodorian realist counterparts TEMPERATURE and BACHELOR. And even if relevance theorists argue that "most occasional representations of a property (or an object, event or state) do not stabilise into a concept" (Sperber & Wilson, 1998, p. 198), they would still have to accept that at least some ad hoc meanings, which are very often used in everyday

[7] For the arguments here, see Sperber & Wilson (1998) and Carston (2012) respectively.

communication, like TEMPERATURE* and BACHELOR* (possibly also BACHELOR**, a stereotypical bachelor who is untidy, or even BACHELOR***, a man who seeks ephemeral relationships and so on and so forth), eventually get to be stored in the mental lexicon; and this time these seemingly ad hoc concepts would effectively be 'semantic', that is, decodable rather than inferred.

This *rampant encoded polysemy*, as Vicente and Martínez Manrique (2010) have aptly called it, would be problematic from a philosophical viewpoint, as it effectively violates Grice's Modified Occam's Razor. And although Carston has noted that such a proliferation of word senses in the lexicon would not necessarily be problematic for RT, since "within a theory of utterance interpretation conceived as a matter of on-line cognitive processes, it might well be more economical to retrieve a clutch of stored senses and choose among them, than to construct an interpretation out of a single sense and contextual information" (2002, p. 219, fn.50), when combined with the preceding discussion of the encoded meaning of words like 'happy' or 'open', this remark certainly raises questions regarding the relevance theorists' need to postulate single, general concepts that encode the meaning of such words in the first place.

Fodor has based his account of semantics on the presupposition that the content of a natural language sentence or a lexical item is entirely isomorphic to some determinate thought or atomic concept that they correspondingly encode. In this respect, apart from the few cases of homonymy, as in the two distinct meanings of the word 'bank', a lexical item carries a single meaning that is referentially derived. That is largely because of the issues that he has sought to address in the first place; issues for which context-sensitivity has traditionally been thought of as problematic, such as compositionality, the assignment of satisfaction conditions to semantically evaluable expressions, intentional explanation and so on and so forth. But this isomorphism does not work when psychological considerations enter the picture, a point that has been made by Sperber and Wilson themselves from at least as far back as (1998). When it comes to actual verbal communication, even Fodor agrees that "language is strikingly elliptical and inexplicit about the thoughts it expresses" (2001, p. 11), but this does not compromise his account, since it has little, if anything, to do with the actual processing of linguistic stimuli per se; it is "an account of the metaphysical character of the (primitive) semantic properties and relations" rather than "a specification of the semantic

properties of the expressions in a language" (Fodor, 2008, p. 18, fn.34). In fact, like most philosophers, Fodor has emphasized time and again that confusing psychology with semantics "is a very bad idea" and that "*there is, as a matter of principle, no such thing as a psychological theory of meaning*", since semantics is by definition about "constitutive relations between representations and the world" (Fodor, 2008, p. 88, emphasis in original).

Relevance theorists are in all likelihood equally aware of the more general dichotomy of interest between philosophers and psychologists who study linguistic meaning as they seem to acknowledge that "it is far from obvious that the label 'concept' refers to the same entity for both parties (and very clear that the term 'semantic' does not) so that little conciliatory progress is likely to be made until these differences are mapped out and resolved" (Carston, 2010a, p. 175, fn.8). Even so, they choose to ignore semantics from their research agenda. For instance, in a recent paper, Carston indirectly responded to the criticisms that RT has been receiving regarding what minimalists perceive to be its semantic commitments by noting that discussions concerning semantics are not central to what the theory is all about and suggesting that the label 'radical pragmaticism' fits the theory's orientation much better than 'radical contextualism'. As she argues, "it is us, the users of language, who are sensitive to context, and, as rational communicating/interpreting agents, we are able, by exploiting this sensitivity in each other, to get linguistic expressions to do a lot more than simply express their standing linguistic meaning" (2010b, p. 266). In this way, Carston distinguishes the study of the cognitive processing that underlies linguistic communication from that of semantic content, or 'standing linguistic meaning' as she calls it. Since RT's concern has always been to account for the ways in which the dedicated inferential process enriches and complements the semantic representation of linguistic strings, the argument goes, it should have nothing more to say about the nature of these representations other than that they are structured strings of Fodorian-style atomic concepts, which in turn need considerable contextual enrichment to reach full propositional status.

As we have seen, however, the implementation of Fodorian-style concepts as actual processing units that lexical decoding feeds into the inferential processor during utterance interpretation appears to be creating problems in its own right; from a psychological perspective, RT's inability to escape the literal-first hypothesis is a case in point, while from a philosophical one, the proliferation of word senses creates an uncontrollable system, where it

becomes very difficult to keep track of how many senses a word effectively encodes. Against this background, relevance theorists might need to explore alternative ways of accounting for encoded lexical meaning, since abstract philosophical models of lexical semantics of the type that Fodor offers do not fit the bill, nor are they supposed to. Therefore, a more psychologically-oriented approach to the question of what an expression's encoded linguistic meaning effectively is seems needed and, to this effect, behavioural research from the domain of psycholinguistics would undoubtedly have a pivotal role to play.

4. Experimental prospects

Given the detailed account of lexical pragmatics that relevance theorists have recently developed, it would certainly be interesting in its own right to see the extent to which the experimental research that has challenged the 'literal first' hypothesis with respect to the processing of figurative language could also be applied to the study of the processing of other types of pragmatically enriched lexical interpretations, like narrowings, approximations, category extensions, neologisms and word coinages. As Wilson and Carston note, some preliminary data on examples from the last two categories (Clark & Clark, 1979; Clark & Gerrig, 1983) suggest that they are "no harder to understand than regular uses" (2007, p. 237), but it remains to be seen how fast the interpretation of neologisms and word coinages takes place in comparison to the processing of literal meaning. Turning to cases of lexical narrowing, like the one presented in (1), and approximation, like the one in (2), I think it would be quite counterintuitive to expect any delays in their processing, but again there is, to my knowledge, a complete lack of experimental evidence to support this intuition.

From the current discussion's perspective, experimentation on different varieties of lexical meaning adjustment would be able to give us a clearer idea of how different types of enriched interpretations of a word relate to its literal meaning. If no significant difference is documented, this would strengthen not only the assumption that literal lexical meaning is in no way exceptional but also the need to come up with more psychologically-oriented accounts of the mental lexicon and more essentially its particular contribution during verbal communication. One of the few experimental studies on lexical pragmatics from an RT perspective that has surfaced in the recent years, for example, suggests that certain context-independent properties of a word's meaning

remain activated even after a metaphorical (Rubio-Fernández, 2007) or narrower interpretation (Rubio-Fernández, 2008) have been reached. However, even though this set of experiments has been extensively quoted in the RT literature, no full explanation has been provided about what these properties are and how they relate to the equally context-independent atomic concept that the same word encodes. Foreseeing that a 'core' relevance theorist's response would be that they are part of the information that is attached to the word's conceptual address and thus not part of its content (as per the discussion of Fodorian mental files above), it still is curious why these particular properties are obligatorily activated during spontaneous interpretation, which in itself gives the impression that they are decoded rather than inferred.

It goes without saying of course that the suggested behavioural research would not be without its limitations either. For instance, it can certainly be argued that it is dangerous to rely too much on time-reaction measurements, which most of the relevant experiments have implemented either way, since they cannot, on their own provide any direct evidence for the contention that we actually interpret literal and non-literal language by using the same types of mental processes.[8] To this effect, using more advanced experimental techniques, and also potentially looking into what neurolinguistics would have to say could provide us with much more solid conclusions. Regardless of any such limitation, however, I am convinced that the tension between the philosophical analysis of lexical meaning and its psychological consideration that the RT discussion of ad hoc concept construction seems to have involuntarily revealed can open up new and exciting prospects not only for theoretical analysis, but also for experimental research on the semantics/pragmatics interface, at a time when the field seems to be dealing almost exclusively with tropes, scalar implicatures and presupposition projection.

[8] For a discussion along these lines from a psycholinguistic perspective, see McElree & Nordlie (1999), a paper suggested to me by John Tomlinson Jr.

REFERENCES

Assimakopoulos, S. (2008). *Logical structure and relevance*. University of Edinburgh, PhD dissertation.

Atlas, J.D. (1989). *Philosophy without ambiguity*. Oxford: Oxford University Press.

Atlas, J.D. (2005). *Logic, meaning, and conversation: semantic underdeterminancy, implicature, and their interface*. Oxford: Oxford University Press.

Bach, K. (1994a). Conversational impliciture, *Mind & Language*, *9*, 124-162.

Bach, K. (1994b). Semantic slack: what is said and more, in S. Tsohatzidis (Ed.), *Foundations of speech act theory: philosophical and linguistic perspectives*. London: Routledge, 267-291.

Bach, K. (1999). The semantics-pragmatics distinction: what it is and why it matters, in K. Turner (Ed.), *The semantics/pragmatics interface from different points of view*, Amsterdam: Elsevier, 65-84.

Barsalou, L. (1987). The instability of graded structure: implications for the nature of concepts, in U. Neisser (Ed.) *Concepts and conceptual development: ecological and intellectual factors in categorization*. Cambridge: Cambridge University Press, 101-140.

Barsalou, L. (1992). Frames, concepts, and conceptual fields, in E. Kittay & A. Lehrer (Eds.), *Frames, fields, and contrasts: new essays in semantic and lexical organization*. Hillsdale, NJ: Lawrence Erlbaum Associates, 21-74.

Borg, E. (2004). *Minimal semantics*. Oxford: Oxford University Press.

Burton-Roberts, N. (2007). Varieties of semantics and encoding: negation, narrowing/loosening and numericals. In N. Burton-Roberts (Ed.), *Pragmatics*. Basingstoke: Palgrave Macmillan, 90-114.

Cappelen, H., Lepore E. (2005) *Insensitive semantics: a defense of semantic minimalism and speech act pluralism*. Malden: Blackwell.

Carston, R. (1988). Implicature, explicature, and truth-theoretic semantics. In R. Kempson (Ed.), *Mental representations: the interface between language and reality*. Cambridge: Cambridge University Press, 155-181.

Carston, R. (1999). The semantics/pragmatics distinction: a view from relevance theory. In K. Turner (Ed.), *The semantics/pragmatics interface from different points of view*. Amsterdam: Elsevier, 85-125.

Carston, R. (2002). *Thoughts and utterances: the pragmatics of explicit communication*. Oxford: Blackwell.

Carston, R. (2004). Explicature and semantics. In S. Davis & B.S. Gillon (Eds.), *Semantics: a reader*. Oxford: Oxford University Press, 817-845.

Carston, R. (2008). Linguistic communication and the semantics/pragmatics distinction, in *Synthese*, *165*, 321-345.

Carston, R. (2010a). Lexical pragmatics, ad hoc concepts and metaphor: from a relevance theory perspective, *Italian Journal of Linguistics*, *22*, 153-180.

Carston, R. (2010b). Explicit communication and 'free' pragmatic enrichment. In B. Soria & E. Romero (Eds.), *Explicit communication: essays on Robyn Carston's pragmatics*. Basingstoke: Palgrave Macmillan, 217-287.

Carston, R. (2012). Word meaning and concept expressed, *The Linguistic Review*, 29, 607-623

Clark, E,. Clark, H. (1979). When nouns surface as verbs, *Language*, *55*, 767-811.

Clark, H., Gerrig, R. (1983). Understanding old words with new meanings, *Journal of Verbal Learning and Verbal Behavior*, 22, 591-608.

Fodor, J.A. (2001) 'Language, thought and compositionality', *Mind & Language*, *16*, 1-15.

Fodor, J.A. (2008). *LOT 2: the language of thought revisited*. Oxford: Oxford University Press.

Gibbs, R.W.Jr. (1994). *The poetics of mind: figurative thought, language and understanding*. Cambridge: Cambridge University Press.

Glucksberg, S. (2001). *Understanding figurative language: from metaphors to idioms*. Oxford: Oxford University Press.

Grice, P. (1989). *Studies in the ways of words*. Cambridge, Mass.: Harvard University Press.

Groefsema, M. (2007). Concepts and word meaning in relevance theory. In N. Burton-Roberts (Ed.), *Pragmatics*. Basingstoke: Palgrave Macmillan, 136-157.

Jaszczolt, K.M. (2005). *Default semantics: foundations of a compositional theory of acts of communication*. Oxford: Oxford University Press.

Levinson, S.C. (2000). *Presumptive meanings*. Cambridge, Mass.: The MIT Press.

McElree, B., Nordlie, J. (1999). Literal and figurative interpretations are computed in equal time, *Psychonomic Bulletin & Review*, *6*, 486-494.

Reboul, A. (2008). Review of Noel Burton-Roberts, *Pragmatics*, *Journal of Linguistics*, *44*, 519-524.

Recanati, F. (1989). The pragmatics of what is said, *Mind & Language*, *4*, 295-329.

Recanati, F. (2001). What is said, in *Synthese*, *128*, 75-91.

Recanati, F. (2002). Unarticulated constituents, *Linguistics and Philosophy*, *25*, 299-345.

Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.

Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Oxford University Press.

Rubio-Fernández, P. (2007). Suppression in metaphor interpretation: differences between meaning selection and meaning construction, *Journal of Semantics*, *24*, 345-371.

Rubio-Fernández, P. (2008). Concept narrowing: the role of context-independent information in concepts, *Journal of Semantics*, *25*, 381-409.

Searle, J.R. (1978). Literal meaning, *Erkenntnis*, *13*, 207-224.

Sperber, D., Wilson, D. (1986). *Relevance: communication and cognition*. Oxford: Blackwell. (2nd edition with postface: 1995).

Sperber, D., & Wilson, D. (1998). The mapping between the mental and the public lexicon, in P. Carruthers & J. Boucher (Eds.), *Language and thought: interdisciplinary themes*. Cambridge: Cambridge University Press, 184-200.

Travis, C. (1981). *The true and the false: the domain of the pragmatic*. Amsterdam: John Benjamins.

Travis, C. (1997). Pragmatics. In B. Hale. & C. Wright (Eds.), *A companion to the philosophy of language*. Oxford: Blackwell, 87-107.

Vicente, A., Martínez-Manrique, F. (2010). On relevance theory's atomistic commitments. In B. Soria & E. Romero (Eds.), *Explicit communication: essays on Robyn Carston's pragmatics*. Basingstoke: Palgrave Macmillan, 42-48.

Vicente, B. (2005). Meaning in relevance theory and the semantics/pragmatics distinction. In S. Coulson & B. Tomaszczyk (Eds.), *The literal and non-literal in language and thought*. Frankfurt: Peter Lang, 179-200.

Wedgwood, D. (2007). Shared assumptions: semantic minimalism and relevance theory, *Journal of Linguistics*, *43*, 647-681.

Wilson, D. (2004). Relevance and lexical pragmatics, *UCL Working Papers in Linguistics*, *16*, 343-360.

Wilson, D. (2011). The conceptual-procedural distinction: past, present and future. In V. Escandell-Vidal, M. Leonetti & A. Ahern (Eds.), *Procedural meaning: problems and perspectives*. Bingley: Emerald, 3-31.

Wilson, D., Carston, R. (2007). A unitary approach to lexical pragmatics: Relevance, inference and ad hoc concepts. In N. Burton-Roberts (Ed.), *Pragmatics*. Basingstoke: Palgrave Macmillan, 230-259.

Wilson, D., Sperber, D. (1981). On Grice's theory of conversation. In P. Werth (Ed.), *Conversation and discourse: structure and interpretation*. London: Croom Helm, 155-178.