

On Relevance Theory's Atomistic Commitments

Agustín Vicente (agusvic@fyl.uva.es) &

Fernando Martínez Manrique (fernan_martinez@yahoo.com)

Introduction

Robyn Carston (2002) has forcefully argued for the thesis of semantic underdeterminacy (henceforth SU), which states that the truth conditional meaning of a sentential utterance cannot be obtained by semantic means alone (i.e. barring indexicals, using fixed lexical meanings plus rules of composition following syntactic structure information). Rather, the truth conditional meaning of a sentence depends to a considerable extent on contextual information. Yet, like other proponents of Relevance Theory (RT henceforth), she endorses atomism of the Fodorian style (Fodor, 1998), which has it that “lexical decoding is a straightforward one-to-one mapping from monomorphemic words to conceptual addresses” (Carston 2002: 141). That is, Carston (and RT in general, see Sperber and Wilson, 1986/1995, 1998) seems committed to Fodor's *disquotational lexicon hypothesis*, whose chief idea is that word meanings are captured in mind by atomic concepts that we can represent by disquoting the word and putting in capitals. Thus ‘dog’ corresponds to DOG, ‘keep’ to KEEP, and so on. This thesis leaves the door open for the possibility that words are related to more than one concept. In fact, RT exploits this possibility by suggesting that the set of concepts is indeed much larger than the set of words. Each word has a corresponding encoded concept that, as we will explain later, can give rise to an indefinite number of different related concepts, constructed ad hoc –in a relevance-constrained manner– for the particular context of utterance.

In this paper we will try to show that SU and DLH are not easy to conciliate. We will argue that the phenomenon of underdeterminacy can be interpreted as a phenomenon of rampant polysemy, which the DLH cannot account for. We also think that the tension existing between SU and the DLH has not passed unnoticed to Carston, and it is the reason why she oscillates between more orthodox versions of RT (2002, Carston and Powell, 2006) and a new one, where meanings are not concepts, but

"concept schemas or pointers to conceptual space" (2002; 360)¹. However, from our point of view, none of these are advisable solutions. Instead, the option that we want to argue for is decompositionalism, that is, the thesis that words express variable complexes of concepts made out of a finite list of typically abstract primitives. A related thesis has already been proposed as an explanation of polysemy by Jackendoff (2002) and Pustejovsky (1995). We want to show that such a style of explanation can also account for the polysemy generated by extralinguistic contextual effects.

The way we will proceed is the following. First we will present and defend Carston's views on semantic underdeterminacy and effability of thought at the sentential level. Then we will try to show how they entail underdeterminacy at the level of words or rules and the rejection of the DLH. After that, we will discuss whether, despite the loss of the DLH, atomism could still be viable: there we will deal with Carston's non-conceptualist view on linguistic meaning (a position she shares with Levinson). It will hopefully emerge from the discussion that decompositionalism is better suited to account for SU and polysemy.

1. Semantic underdeterminacy for sentential utterances: the ineffability of thought

As stated above, the thesis of semantic underdeterminacy claims that the truth conditional meaning of a sentential utterance cannot be obtained by semantic means alone. Our claim is that if SU is a widespread phenomenon (as Recanati, and, especially, Carston claim), then Fodor's conceptual atomism is in trouble to account for lexical concepts. In fact, if there is hope for any atomistic theory, it will have to be without the Disquotational Lexicon Hypothesis. SU is usually stated as the thesis that the propositional (truth-conditional) content of sentential utterances is not determined by semantic values and compositional (non-pragmatic) rules of language alone². It is a

¹ We can advance the reason why RT orthodoxy is not available to Carston. It is true that RT has from its very beginning defended that most natural language sentences are semantically underdetermined: decoding a sentence does not give you a proposition, but a propositional template. However, such underdeterminacy is localized, so to speak, for it is generated by the presence of indexicals, free variables and other known context sensitive expressions, such as gradable adjectives. Carston (2002) departs from this view, holding that underdeterminacy is more general and can affect all referential expressions and all predicative terms. But this departure has a price, as we will try to show and have already claimed: rampant polysemy and the rejection of DLH. Meanings then cannot be concepts, if these have to be atomic. That is why Carston proposes that they are less than conceptual in the final chapter of her book.

² We will assume for the time being, and along with common wisdom, that such rules are determined by syntactic information.

thesis motivated by reflection on a variety of examples, ranging from 'John's car is empty' (Recanati, 1995) to 'the kettle is black' or 'those leaves are green' (see Carston, 2002, Travis, 2000). The truth-conditions of an utterance of 'John's car is empty' depends at least on what relation holds between John and the car, and also on the sense in which such car is said to be empty. An utterance of 'the kettle is black' may be true if the kettle is totally painted in black, but also if it is burned, or dirty, or just partially painted in black. It may be claimed that these utterances have a literal meaning, a meaning that is constant across contexts. However, we take it that this kind of meaning, that would be of a very general nature, cannot be said to capture the truth-conditions of the utterance. (In the case of 'John's car is empty', this meaning could be "a thing that is said to be a car (real car, toy car, or whatever) that is in some close relationship to someone presumably called John is relevantly empty"). We would not want to discard that this meaning is not even propositional.

Now, defenders of SU use these and other examples to motivate their thesis. However, it is possible to object that even if this were true about many sentential utterances, it does not necessarily apply to most sentences that one can *potentially* utter. Imagine that for any underdetermined sentential utterance it is possible to find another sentence that explicitly, or literally (i.e., without any contextual aid) expresses the same proposition that the former expresses aided by contextual information. That is, imagine that the underdetermined 'the kettle is black' is translatable into 'the kettle is burned' and that this latter utterance has the propositional content of the former as its literal meaning. Then, it seems, we are in a position to claim that 'kettle', 'burned', etc. have themselves a literal meaning, namely, KETTLE, BURNED, etc. Words would not always express the same concepts (perhaps they would rarely do), but the DLH would stand: the meaning of a word could be an atomic concept.

We want to argue, following Carston, against the existence of such eternal sentences, but first we want to point out that the recourse to eternal sentences may not be particularly helpful for the atomist defending the DLH. To take a new example, think of 'the beach is safe' (Fauconnier and Turner, 2002). This may mean (at least) that the beach is protected from certain dangers, or that the beach is not dangerous in known ways. That is, 'the beach is safe' is semantically underdetermined, and its different meanings can be captured by putative eternal sentences (obviously, they are not so, but let us assume they are for the sake of the argument) such as 'the beach is protected from harm' and 'the beach is not harmful'. What this implies is that SAFE is not one of our

atomic concepts: it is a variable complex, a complex formed by at least the concept HARM. As semantic underdeterminacy, although perhaps not completely general, is general enough, we would then have that the stock of primitive concepts is considerably smaller than the lexicon of English. So perhaps the atomist would be ill advised to hold onto eternal sentences, given that it looks like a sure road to decompositionalism.

Now, the discussion about eternal sentences is long and sustained (see Quine, 1960, Katz 1978, Wettstein, 1979, Recanati, 1994, Carston, 2002). One of the most recent defences of generalized underdeterminacy is due to Carston (2002), where she argues against the following “principle of effability”:

“For every statement that can be made using a context-sensitive sentence in a given context, there is an eternal sentence that can be used to make the same statement in any context” (p. 34)³.

What she goes on to show is, first, following Wettstein (1979) that there are no univocal translations of indexicalized, or, in general, context-sensitive sentences to putative eternal sentences which replace all context sensitive expressions with uniquely denoting descriptions (‘She left in a hurry’ can be translated into ‘the woman who spoke to Tony Blair at t_1 left in a hurry’ just as well as into ‘the woman in the red velvet dress who was in the Islington town Hall between t_1 and t_2 left in a hurry’). Now, given what the principle of effability says, this is not a proof against it: if one reads “there is an eternal sentence” as “there is at least an eternal sentence” (and we find that reading plausible), then it seems that Wettstein-style responses concede effability.

However, Carston provides a second reason to question effability, that we find more convincing. What she argues is that there are no such things as eternal reference and eternal predication, that is, that all referential and predicative terms are open to contextual variation: any referential term can be used to refer to an actual object, a possible object, or a belief-word-object, while one could say that predicative terms are always interpreted according to a “certain understanding”, as Travis, another defender

³ Another recent, more moderate version of this claim can be found in Bach (2005: 28): “for every sentence we do utter, there is a more elaborate, qualified version we could utter that would make what we mean more explicit”. The ‘elaborate version’ is closer to what is literally meant, but it is unclear whether he conceives of it as a version that is always perfectible (in which case there would not be eternal sentences as such).

of ineffability would put it (see Travis, 2000)⁴. It seems to us that the former point can be illustrated by means of Culicover and Jackendoff's Mme. Tussaud's examples. C & J hold that 'Ringo is the Beatle that I like the most' may be asserted about the actual Ringo or the statue of Ringo. Now, think about the more precise formulation 'The statue of Ringo is the one I like the most'. This should clarify the possible ambiguity, but does it? Well, it clarifies *that* ambiguity, but this new sentence happens to be ambiguous itself, for we could be speaking either about the statue of Ringo or about a picture of it (we took pictures of all the statues and the picture of Ringo's statue happens to be the best). Culicover and Jackendoff sum up the lessons of these examples in the following "statue rule" (p. 371):

"The syntactic structure NP may correspond to the semantic/conceptual structure PHYSICAL REPRESENTATION (X) where X is the ordinary interpretation of NP".

This rule is by itself enough to jeopardize effability. But it is probably just one source of variation among many others (think of a belief-world rule that enables a NP to refer to a mental representation of its usual reference)⁵. And, as already noted, there are reasons to think that predicative terms are also importantly context sensitive: 'the kettle is black' can be translated into 'the kettle is burned', but this latter sentence is as underdetermined as the first was, although not in the same ways. (On this see Travis, 2000). As things stand, we take it that the burden of proof should be on the defender of eternal sentences.

2. Words and Concepts

We claim that the SU of sentential utterances puts in jeopardy the possibility that lexical items (words, for short) have atomic literal meanings themselves. In the mentalistic framework we are assuming, the literal meaning of a word will be understood as the context-independent mental representation that it *encodes*. (There will typically be other concepts that the word *expresses*).

⁴ We think that the underdeterminacy present in "Travis cases", such as 'the ink is blue', is better accounted for by an indeterminacy in the rules of composition (see below). However, this is a minor point in the present context, where we are dealing with Carston's views on effability.

⁵ Chomsky's (2000) discussion of what 'London' can stand for also shows that the semantic contribution of at least some proper names is highly underdetermined: 'London has moved to the South Bank' will probably mean that the fun in London is to be found on the South Bank, but it could also mean, in another context, that the buildings of the city have been physically moved to the South Bank (e.g., due to some coming earthquake on the North of the Thames in a sci-fi movie).

Some of the examples about semantic underdeterminacy presented above apply in a straightforward manner to words. For instance, the claim is that ‘London has moved to the South Bank’ is underdetermined partly because ‘London’ is underdetermined itself. But doubts about the generalizability of this case arise more naturally for words. In what follows, we are going to examine several possibilities in which the idea that words have literal meaning is fleshed out in a way that favours conceptual atomism, i.e., possibilities in which the literal meaning of a word is understood as an *encoded atom* that corresponds one-to-one to the word. Our rejection of such possibilities will rely on a simple argument: if words’ literal meanings are understood as encoded atoms, then it will normally be possible to construct eternal sentences out of them; but since eternal sentences are suspect, so are literal meanings for words.

To begin with, it could be tempting to say that what SU shows is that the semantic value of the whole (the sentence) cannot be fully determined from the semantic value of the parts (the words) plus rules of composition, not that the latter have an underdetermined value as well. In fact, Recanati himself has defended an *accessibility-based* model of utterance processing that fits nicely with the idea of there being literal word meanings (Recanati, 1995, 2003). In his model, the interpreter of an utterance recovers initially the literal values of constituents, but not a literal interpretation of the *whole*. Instead, the interpreter can reach first a nonliteral interpretation of the utterance, provided that the elements of this interpretation are associatively derived, by spread of activation, from the literal constituents that were initially accessed.

This model might be attuned to make it compatible with conceptual atomism in the following way. First, in hearing ‘the kettle is black’ the interpreter accesses the encoded conceptual atoms corresponding to the words (e.g., BLACK). Then, by a contextually constrained process of spread of activation, another concept (say, BURNED, or perhaps a compound of concepts) is reached. Finally, all the concepts ultimately operative are subject to composition to obtain the thought that corresponds to the interpretation of the sentence. However, this atomistic model can run into trouble if it is conceded that (the majority of) the operative atomic concepts have a word that encodes them. For instance, if BURNED were the concept encoded by ‘burned’. This entails that, in most cases, there would be a natural language sentence that translates literally the thought reached by the interpreter, i.e., a natural language sentence formed by the words that encode each of the conceptual atoms. However, this amounts to defending the

existence of eternal sentences. So, if there are reasons to reject them, there are reasons to distrust the model under consideration.

However, this does not make the eventual rejection of the DLH mandatory. For, first, it is possible to say that words are related to concepts in a one-to-one way, even though words can be used to express an open range of concepts, giving thus rise to underdeterminacy effects. The inexistence of literal meanings for sentential utterances could be explained by holding that these various concepts that words can express are not lexicalised (that is, if an utterance of, say, ‘angel’ does not express the coded concept ANGEL, but a concept related to being kind and good, this second concept will not be lexicalised: thus, we avoid being finally committed to effability in the sense mentioned above). Second, it is also possible to argue that the semantic underdeterminacy of the wholes is not due to the underdeterminacy of at least some of its parts, but to the rules of composition. Let us examine these two positions in turn.

Ad hoc concepts

A way to flesh out the first position comes from a variant of Relevance Theory that we will call (in order to distinguish it from the two-level account we will discuss later) ‘conceptual Relevance Theory’ (see Sperber and Wilson, 1998; Wilson, 2003, Carston and Powell, 2006). Its key idea is that words *encode* concepts in a one-to-one way, even though they can *express* an open range of them. Explaining how we go from the encoded concepts to the expressed ones is the task of the new field of “lexical pragmatics” (see Wilson, 2003). Basically, lexical pragmatics has to explain core processes such as narrowing and loose talk or widening, by which we modulate the extension of a given encoded concept. Sperber and Wilson (1998) offer the following piece of a dialogue as an example:

Peter: Do you want to go to the cinema?

Mary: I’m tired.

It is clear that Mary wants Peter to understand that she does not want to go to the cinema. However, in order to do so, Peter has to understand first that Mary is tired in a particular way –tired enough not to go to the cinema. So, Sperber and Wilson (1998: 194) conclude that “Mary thus communicates a notion more specific than the one encoded by the English word ‘tired’. This notion is not lexicalised in English”.

According to conceptual RT, modulated notions are *ad hoc* concepts built on-line from encyclopaedic knowledge in order to meet the requirements of communicative relevance. However, these concepts would be just as atomic as encoded concepts are. So rather than a construction of concepts what we have may be a search and activation of concepts that fit the demands of relevance. There are, however, some problems for this account. First notice that one must face a problem of ineffability, in the sense that it is not possible to give a verbal expression to the concept corresponding to an uttered word (except by uttering it again in the same context): it may be just any one inside of the range of concepts that the word can express. We prefer not to make much of this problem because ineffability, as we see it, is a difficulty for any theory of concepts that embraces underdeterminacy⁶. Where we see some problems, however, is in the role assigned to *encoded* concepts, and their relation to *ad hoc* concepts, which, remember, are also atomic, not just complexes of encoded concepts.

A first problem is the following: RT's preferred explanation is that *ad hoc* concepts are obtained inferentially from encoded concepts. According to this account, the necessary steps for inference to take place would be, first, decodification of words onto their encoded concepts, and, crucially, their composition into a representation. Some of these representations would be propositional templates: this would be the case of representations corresponding to sentences with indexicals, free variables, gradable adjectives, etc. However, some of them would be fully propositional: the conceptual representation corresponding to, say, 'some leaves are green' would putatively be SOME LEAVES ARE GREEN. But such a representation would in effect constitute the truth-conditional non-contextual meaning of the sentence, and this is exactly what SU denies.

Put in other words: That sentences do not encode propositions, but propositional templates, i.e., that there is no literal propositional meaning at the level of sentences, is one of the most basic tenets of RT. It is precisely this feature of RT what distinguishes it from minimalist options such as Emma Borg's (2004), where the pragmatics module (in this case, the central processor) takes as its input a propositionally interpreted logical

⁶ Fodor (1998) poses a problem of ineffability to decompositionalism of the sort we want to defend. Basically, he suggests that a putative primitive such as CAUSE is as polysemous as the complex KEEP (=CAUSE [ENDURANCE OF STATE X]), unless the concept CAUSE is not the concept encoded by the word 'cause'. But if this is so, he argues, the concept CAUSE is ineffable, something which, according to his view, is problematic on its own. However, if ineffability is a consequence of underdeterminacy, this would not be a problem for the decompositionalist, but to any contextualist, perhaps including Fodor (2001).

form⁷. Yet, it looks as though if you have a full composition of concepts you must have a proposition, and full compositions are in this model all over the board. This being so, barring indexical expressions and the saturation of free variables, literalism (or, in any case, the DLH) at the level of the lexicon brings in its wake a kind of minimalism, i.e., the existence of propositions in a first stage of processing. This, besides, resurrecting effability, is problematic when we take into account Travis cases such as ‘the ink is blue’ or ‘the leaves are green’. Which one of the several propositions expressible by utterances of these sentences is *the* encoded one?

Second, commitment to *ad hoc* concepts has trouble with what we can call “semantic polysemy” (‘encoded polysemy’ would be the way to put it in RT terms). Think of ‘window’ in ‘he crawled through the window’, ‘the window is broken’, ‘the window is rotten’, and ‘the window is the most important part of a bedroom’ (Pustejovsky, 1995). There are two atomistic possibilities to explain this variability in the meaning of ‘window’: either there is just one concept WINDOW which is ‘window’s literal meaning, or there are at least four atomic concepts corresponding to it: WINDOW*, WINDOW**, and so on. Now, RT cannot explain how we would go from WINDOW to any modulation of it, since it is a modulation that does not depend on any pragmatic processing (that’s why we called it ‘semantic polysemy’). So RT would rather say that ‘window’ encodes not one, but various concepts (hence the label ‘encoded polysemy’), thus departing from the DLH. However, it would still be difficult for RT to explain why ‘the window is broken’ activates WINDOW**, instead of any of the other three concepts. Their defenders could try to say that this specific activation is due to the activation of BREAK, such that there is a sort of co-composition in the decodification process. But then, how does this co-composition happens? It is easy to explain co-composition if you are a decompositionalist, for you can say that BREAK activates

⁷ The question of the psychological reality of minimal propositions can be approached by distinguishing at least three kinds of positions: minimalism first, minimalism too, and minimalism if. Minimalism first has it that minimal propositions are necessary steps in linguistic processing. Minimalism too claims that minimal propositions are *also* entertained, although they are not necessary steps in the processing. Minimalism if would hold that minimal propositions can be entertained and so be psychologically real, only that not always and not necessarily (just if some contextually relevant information is absent). From the two minimalist theories today fashionable, we take it that Cappelen and Lepore’s is a minimal too theory, while Borg’s is a minimal first. If we are not misfiring, RT would have to claim that (sometimes) there are minimal first propositions too.

GLASS (or a glass-related concept), as a part of the complex concept WINDOW. But we do not see how the story would go for an atomist⁸.

Now, these may be seen as problems derived from the inferentialist commitments of RT. A non-inferentialist model, such as Recanati's, seems free of them. In this model the interpreter of an utterance recovers initially the literal values of constituents, but not a composition of them. What the hearer composes are not literal meanings or encoded concepts but concepts already modified by linguistic and extralinguistic contextual information. Thus, it is possible to avoid not only the commitment to non-contextual truth-conditions (minimal propositions), but also the problem derived from semantic polysemy.

The trouble with this model is what role the first concepts do really play. If they are intermediate stations, that do not become part of the interpretation, then it is unclear that they are playing any semantic role after all. On the other hand, it is not easy to get a clear idea as to how modulation processes go either in this model or in conceptual RT. According to RT, lexical entries are nodes in memory composed by three kinds of constituents: atomic concepts, encyclopaedic, or world, knowledge, and a logical entry which captures its analytic entailments. On the latest versions of the theory, such as Carston's (2002) or Wilson's (2003), lexical pragmatic effects (loosening and narrowing) are accounted for by means of the second of these constituents: encyclopaedic knowledge. If, speaking of a friend, we say 'Anne is an angel' we expect the hearer to search in her encyclopaedic information associated to 'angel', and widen up its extension so as to meet her expectations of relevance. This will be done when she takes into account some particular properties angels are supposed to have, say, being kind, being good and a number of others, and excludes others, such as having wings and having no well-established sexual identity. Then, according to RT, the hearer comes to entertain the concept of ANGEL*, a new, *ad hoc*, concept with a wider extension and different encyclopaedic knowledge associated.

Now, we find this picture appealing and promising, safe from one point, related to our concerns here: how are *ad hoc*, or, in general, modulated atomic concepts, construed or reached? It seems to us that the explanation would be simpler and more intelligible if *ad hoc* concepts were complexes made up of atomic concepts. The new node 'angel*'⁹,

⁸ For the full story of co-compositions, see Pustejovsky (1995)..

⁹ This would not be a word of our language, but a representation in the LOT.

from our point of view, seems exhausted by a composition of the entries 'kind', 'good', and several others. Thus, the concept ANGEL*, i.e. the extension of the entry, seems to be given by the composition of the extensions of KIND, GOOD and the other concepts, while it looks like its encyclopaedic information equally consists of the similar composition of the encyclopaedic informations pertaining to 'kind', 'good', etc. Of course this turns ANGEL* into an effable, lexicalisable concept. It may be difficult, perhaps impossible, to give an adequate characterization of the concept. Surely it is not exhausted by 'kind and good'. But, as a matter of metaphysical fact, we can say that there is such a characterization since the concept ANGEL* and the representation 'angel*' are made up, respectively, of encoded lexicalised concepts and of lexical entries. In a nutshell, we think that the best way to make sense of modulations in general is by explaining them as concept compositions. Leaving at one side the problem of the implications of this for the SU thesis (it contradicts the thesis of ineffability), the point is that the atomist owes an explanation of modulation processes consistent with her atomism¹⁰.

Given the above –for us, simpler- understanding of *ad hoc* concepts, it is possible to speculate about the reasons that relevance theorists may have for defending their atomistic nature. There are several by now classical reasons for atomism in general, most of them due to Fodor (see Fodor, 1998). Very roughly and compressed, they can be summed up in the idea that only atomism can explain some features of concepts such as (i) the inexistence of definitions for them, (ii) their intersubjective shareability, (iii) the compositionality of thought, and (iv) the denotational determination of natural-kind concepts. We tend to think that RT's commitment to atomism is grounded on these classical reasons put forward by Fodor. However, these, if at all, are good reasons for having *encoded* concepts atomic. *Ad hoc* concepts are not supposed to be intersubjectively shared or even shareable. But even if they were, this would not demand that they be atomic, as long as, as we propose, they could be made up of atomic concepts. Second, there is no problem of definitions for *ad hoc* concepts. As we said, it is possible that we cannot come up with a definition of ANGEL*, but this would not mean that there is such a definition, which, in any case, would be context-dependent. Third, if ANGEL* is composed by atomic concepts, there would be no problem of

¹⁰ It is to be noted that we might have a problem of massive conceptual innatism lurking behind this picture, given that we need an atomic concept C* for every possible sense in which a word is uttered.

compositionality: the thought ANNE IS AN ANGEL* would be given by the composition of the singular concept ANNE, and the concepts of BEING KIND, BEING GOOD, etc. Finally, the problems of ignorance and error, which haunt decompositionalist accounts of concepts when applied to natural kinds, seem misplaced when the topic is *ad hoc* concepts. Could we be wrong about the extension of an *ad hoc* category, construed on-line for the purposes of understanding? We cannot see how. The extension of LION does not depend on its stereotype or a definition such as FEROCIOUS BIG CAT-LIKE PREDATOR, but LION* (the *ad hoc* concept triggered by 'my man is a real lion') is not a natural kind concept, and so its extension is under our control, so to speak. For instance, it could be well determined by the concepts DETERMINED and FEROCIOUS.

What we would want to conclude, then, is that there are no deep reasons to consider that *ad hoc* concepts must be atomic, just like encoded concepts. And if there are no such reasons, then our view that they have a complex structure should be the default hypothesis, given that the resulting picture is more parsimonious. The problem, as explained above, is that under this view *ad hoc* concepts are in principle lexicalisable, which compromises the idea of ineffability.

3. Rules and Concepts

So far we have argued that conceptual RT may have problems in conciling SU with atomism. We have first shown some concerns about how RT, given its commitment to a first-stage composition, would avoid the existence of "minimal first" propositions, and deal with encoded polysemy. Then we have examined the recourse to *ad hoc* or modulated concepts. To repeat our criticism, common to inferentialist and associationist readings of the overall schema, we think that, taken that schema as granted, the default view on what modulation consists in should be that it is a process of composition of encoded concepts, which turns modulated meanings into encodable or lexicalisable. If this is generally so, then there would be minimal propositions, i.e. propositions that are encoded by sentences, which retain their meaning across contexts of utterances. This, we think, falsifies the idea that there are no eternal sentences. As we also think that Carston's arguments and Travis-like examples are good against the existence of such sentences, our momentary conclusion is that the DLH cannot stand.

However, to end the discussion of the argument we have to consider, as we advanced above, another possibility: whether the semantic underdeterminacy of sentences may be due not to their constituents but to rules of composition. It might seem plausible to say that, e.g., ‘the ink is blue’ (Travis, 2000) is underdetermined (the ink may be blue in its writing, or it may be blue in its external appearance, though it writes black) because it is not specified how ‘blue’ is supposed to modify ‘ink’. That is, there is no underdeterminacy in the meanings of ‘blue’ or ‘ink’: ‘blue’ means BLUE and ‘ink’ means INK. What makes the sentence underdetermined is the indeterminacy of the composition rule.

We doubt that this kind of response would be popular. Orthodoxy has it that composition rules follow syntactic information, and the syntactic structure of ‘blue ink’, or of ‘the ink is blue’ is not ambiguous, so the composition rule cannot be either. However, one might want to depart from orthodoxy here, and claim that semantic composition rules do not track syntactic structure. Perhaps there are properly semantic rules, so that adj + noun or noun + noun constructions are underdetermined because it is not determined which of the various possible semantic rules we have to follow in a given case. Sometimes an adjective such as ‘blue’ modifies its head by being applied to its referent considered as a stuff of a certain kind, while other times its modification is applied to what its referent typically does. Now, the problem with this response is that it probably hides a trap to the atomist. For the best explanation for the variation in the composition rules is that meanings are complexes, such that, e.g., an adjective may be applied to different parts of the complex. That is, if composition rules worked as suggested, then it seems that concepts would have to have a certain inner structure. There must be information within the concept about what the things that fall under it are for, or even what kind of thing they are: if PEN is decomposable into at least PHYSICAL OBJECT, and, say, USED FOR WRITING (though this is a rough proposal: WRITING should be decomposable in some other basic concepts), then it is possible to explain that there are at least two kind of rules of composition that can be applied to RED PEN. In contrast, it seems difficult to explain how come the atomic RED can be applied in various ways to the atomic PEN: if RED applies to PEN, it seems all that you can possibly have is the complex RED PEN, whose denotation is the intersection of red things and pens.

To sum up, it is *prima facie* open to the atomist to argue that the semantic underdeterminacy of sentential utterances does not establish the semantic

underdeterminacy of some of its parts, since sentence underdeterminacy may be caused by underdeterminacy of rules. However, if this path is taken, it will have to be at the price of conceding that concepts have some kind of internal structure all the same. Furthermore, it will have to be conceded that, lacking contextual information, it is underdetermined which parts of the decomposition have to be activated: which is to say that, lacking contextual information, what the word means is underdetermined. We tend to think that, in effect, the phenomenon of sentential SU is not homogeneous. Some cases are due properly to the underdeterminacy of constituents and some others to an underdeterminacy induced by the inespecificity of the rules of composition. In the first case, we think it proper to defend a decompositionist account, according to which mental representations of word meanings consist of a cluster of primitive concepts with perhaps a basic, more or less stable core, and a number of variable components that can be adjoined or removed depending on particular contexts of utterance. In the second, we find it reasonable to regard concepts as structured entities, and account for differences in meaning as differences in what parts of the structure are brought to focus by the composition rules (a proposal much in the line of the co-composition processes defended by Pustejovsky, 1995). What is important to note is that in either case, the DLH turns to be false: concepts are not atomic, and words are not univocally related to concepts.

To close this section, we want to remark that even if one thinks that semantic underdeterminacy is not so widespread as many pragmaticians hold, the phenomenon is still general enough so as to undermine the plausibility of Fodorian conceptual atomism, which actually predicts widespread literalism. Nevertheless, this variety of atomism relies on the existence of a single level, the conceptual level, to capture the meaning of words. If we differentiate semantic from conceptual affairs, it is possible to put forward an atomistic model without the burden of DLH. The idea, in a nutshell, is to posit the existence of semantic entities that encode word meanings, while leaving the solution of semantic underdeterminacies in the hands of the conceptual, pragmatically-driven device. It is to this idea that we turn now.

4. Two-Level Theories

A two-level theory accounts for the relation between word meanings and concepts by means of two different processing levels: semantic and conceptual. Levinson (1997,

2003) has presented a number of different reasons in support of this distinction. The key idea is that language is necessarily too general to be able to encode the specific thoughts we entertain. In consequence, representations of linguistic meaning have to be different from representations that actually appear in thought, i.e., semantic and conceptual representations must be kept in separate mental levels. There are partial mappings from one level onto the other, but the respective representations are not isomorphic. The distinction can be used to safeguard linguistic diversity while avoiding the pitfall of unrestricted linguistic relativity: different languages might differ on the semantic representations they encode, while conceptual representations (the stuff on which, according to Levinson, “serious thinking” operates) are constructed out of (roughly) universal elements. So one could say that semantic representations invoke thoughts, rather than encoding them.

Carston’s (2002) defence of Relevance Theory can be regarded as another version of a two-level theory. We mentioned above that it is possible to distinguish two different trends in RT, a conceptual RT and a two-level RT. Both trends are present in Carston’s work (see e.g. Carston and Powell, 2006 for the first), but here we will focus on the second of these.

According to Carston, then, word meanings are not concepts, but rather “concept schemas or pointers to conceptual space” (see Carston, 2002: 360). There is an obvious motivation for holding such a position, and it is her determinate defence of massive underdeterminacy. It seems that, on her account, logical forms underspecify the meaning of utterances in all possible ways, and this includes, as we said in the beginning, both referential and predicative terms. If this is so, then probably what a given word encodes is just a clue to search for a concept depending on contextual information, much in the way that indexical terms do. Therefore, in this regard Carston would be sharing Levinson’s rationale for his distinction of levels: to repeat, that language is of necessity too general to encode our thoughts.

Now, it is possible to speculate whether there may not be some further reason for Carston and Levinson (although see also Powell, 2000) to defend a two-level view. In particular, it can be thought that, given the underdeterminacy thesis, the compositionality of the semantics of language demands such a distinction. The underdeterminacy thesis has it that the decoding of a natural language sentence typically either does not convey a full proposition or conveys a proposition different

from the intended one –which can even be taken to be “the meaning of the sentence”. Now –so the argument would go– our understanding of a natural language is productive and systematic, in a way that can be explained only by assuming that it has a compositional semantics. And, given the underdeterminacy thesis, this means that the units that are composed when a sentence is decoded cannot be concepts, for a full composition of concepts is a thought, i.e., a truth-valuable whole. So, there is good reason to distinguish semantic and conceptual levels in linguistic processing: to repeat, language is compositional in its semantics, but natural language sentences do not express thoughts literally (by decoding). So the semantic units that are composed cannot be concepts.

Now, this argument seemingly arises not from a compositionality principle but from a demand imposed by the view that all linguistic processing other than decoding is inferential in nature. In Relevance Theory the only heuristics involved in language understanding is the principle of relevance. According to this principle, understanding an utterance relies on inferential processes that take “wholes” as input and then produce more relevant “wholes” as output. Now, the phenomenon of underdeterminacy shows that the input cannot be a proposition, i.e., something that corresponds to a fully truth-conditional thought. So it has to be a “whole” obtained compositionally by different means. In other words, there must be a different, non-conceptual level where semantic composition takes place, so as to provide the starting “whole” for inferential processes to begin their job. However, if one drops the demand that all pragmatic processing is inferential then one is not necessarily committed to the mental reification of such a distinct semantic level. In other words, whatever semantic elements are involved in understanding an utterance, they can make their contribution together with pragmatic constraints, and not prior to them.

Yet, one might insist that the starting point of our linguistic understanding are non-conceptual semantic units, even if we do not compose them at early stages of processing. In our view, there are more general grounds to resist this distinction of levels. The problem, in a nutshell, is to offer a reasonable characterization of the distinction, saying (a) what the semantic elements are, and (b) how they give rise to concepts. With respect to the first question, there are two different kinds of answers: either they are regarded as representational structures, or they are not. For instance,

Levinson opts for the representational approach while Carston seems to hesitate between both answers in saying that they are “concept schemas” (arguably a representational solution) or “pointers to conceptual space” (arguably a non-representational one).

If elements at the alleged semantic level are not representations, we arrive at a position where words just initiate a pragmatically-guided activation of conceptual units that renders a different concept for each different context of utterance. It may well be accepted that in the process of interpretation there are pointers or relay stations that activate different regions of the conceptual space. What is thoroughly rejected is that there is any intermediate interpretation constructed *out of those pointers*. Hence, if the two-level theorist is to avoid this kind of picture, and the consequence of being unable to explain the (relative) stability of meaning, she ought to insist on a representational reading of the semantic level.

In fact, typical discussions regarding the semantic/conceptual divide tend to pose the question in representationalist terms. The issue is which elements of the mental lexicon need to be represented in a separate semantic lexicon, and which ones belong to a (more general) conceptual structure. Offering a detailed analysis of the problem goes beyond the purposes of this paper (see Murphy, 2003, ch. 3, for an up-to-date review of positions on that issue). Still, there are a few considerations that undermine the possibility of using a distinction between levels to sustain atomism.

The chief problem is how to account for the relation between semantic and conceptual representations. To begin with, one has to do so in a way that avoids duplication of elements in the mind. Levinson (2003) seems to commit this mistake when he talks about a certain kind of concepts that are closer to the semantic realm: those would be the concepts onto which semantic representations immediately map. But those concepts serve the same purposes as the semantic representations were meant to serve, so it is unclear why the latter should be posited in the first place. So one must suppose that semantic units are different from conceptual elements in the way they function, and this is what a characterization of semantic elements in terms of concept templates, pro-concepts and the like is meant to convey. But the characterization is still too vague, and it is not easy to see how to sharpen it. A possible way could be to hold that semantic units are characters, functions, or specifiable rules connecting words to concepts (and perhaps these characters are non-conceptual, which is to concede much). It is true that these accounts are usually

semantic-conservative in a way that Carston and Levinson are definitely not, but maybe one could work out a version of their strategy of indexicalization palatable to them (or rather, atomists, in general). If this were done, then there would be a way to say not only what semantic units are –they are characters– but also how they give rise to concepts –just in the way characters give rise to semantic values. Unfortunately, it is doubtful that such a strategy could do the work the atomist wants it to do.

Remember that the atomist has to explain, among other things, how words like ‘Sampras’, ‘small’, ‘tired’ or ‘angel’ may stand for a variety of concepts. That is, what she has to explain is how (most) words in a language turn out to be polysemous. The present attempt consists in saying that they contain rules of use much like the characters of indexical expressions that point to concepts. The question is: what kind of character could link the term ‘Sampras’ to the concept, say, of COMPLETE PLAYER THAT DOMINATES AN AGE?¹¹ In our view, the complexity of the required link makes it unlikely that characters could be simple, unstructured entities.

Conclusion

We can now summarise what we have been defending. We have begun by trying to support the claim of those who, like Travis, Carston and others, hold that most, if not all, sentential utterances are semantically underdetermined. We have suggested that this claim implies that thoughts, or propositions, are not in general effable, that is, encodable by sentential utterances alone (this, we think is what especially Travis-cases make manifest). Then, we have argued, especially *contra* conceptual RT, that the responsibility of the underdeterminacy of sentential utterances must rest either on their constituent parts or on their composition rules. In either case, we have claimed, the DLH is false: meanings cannot be atomic concepts, and they cannot be related one-to-one to words. Last, we have examined the atomist position that distinguishes two levels in linguistic meaning, the semantic and the conceptual, concluding that it is a proposal that, at the present stage, seems unwarranted and unclear in important respects. From all this we would like to conclude finally that Fodor’s atomism should

¹¹ Notice that the fact that one can use a proper noun with a conceptual value shows that simple links cannot do the trick. If one limits oneself to fill the reference of the name one obtains an identity claim that is manifestly false.

be rejected, especially by those committed to claims of underdeterminacy or ineffability. We think decompositionism is in a better position to account for such putative facts, so that it should be reconsidered by students of linguistic meaning. Unfortunately, we have to postpone for a future paper a proper defence of decompositionism, necessary to support this last claim.

References

- Bach, K. 2005: 'Content *ex Machina*', in Z. G. Szabó (ed) *Semantics versus Pragmatics*, Oxford: Oxford University Press.
- Borg, E. 2004: *Minimal Semantics*. Oxford: Oxford University Press.
- Cappelen, H. and Lepore, E. 2005: *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.
- Carston, R. 2002: *Thoughts and Utterances*. London: Blackwell.
- Carston, R. and Powell, G. 2006: Relevance Theory –New Directions and Developments, in Lepore, E. and Smith, B.C., eds. *The Oxford Handbook of Philosophy of Language*. Oxford University Press.
- Chomsky, N. 2000: *New Horizons in the Study of Language and Mind*, Cambridge: Cambridge University Press.
- Culicover, P. and Jackendoff, R. 2004: *Simpler Syntax*, Oxford University Press.
- Fauconnier, G. and Turner, M. 2002: *The Way We Think*. New York: Basic Books.
- Fodor, J. 1998: *Concepts*. New York: Oxford University Press.
- Fodor, J. 2001: 'Language, Thought and Compositionality', *Mind and Language* 16, pp. 1-15.
- Jackendoff, R. 2002: *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford: Oxford University Press.
- Katz, J.J. 1978: "Effability and Translation", in Guenther, F. and Guenther-Reutter, M., eds. *Meaning and Translation: Philosophical and Linguistic Approaches*, 191-234. London: Duckworth.
- Levinson, S. 1997: 'From Outer to Inner Space: Linguistic Categories and Non-linguistic Thinking', in E. Nuyts and J. Pederson (eds) *Language and Conceptualization*, Cambridge: Cambridge University Press, pp. 13-45.

- Levinson, S. 2003: 'Language and Mind: Let's Get the Issues Straight', in D. Gentner and S. Goldin-Meadow (eds) *Language in Mind*, Cambridge: MIT Press, pp. 25-45.
- Murphy, M. L. 2003: *Semantic Relations and the Lexicon*. Cambridge: Cambridge University Press.
- Powell, G. 2000: "Compositionality, Innocence and the Interpretation of NPs". *UCL Working Papers in Linguistics* 12: 123-44.
- Pustejovsky, J. 1995: *The Generative Lexicon*. Cambridge, MA: MIT Press.
- Quine, W.V.O. 1960: *Word and Object*. Cambridge, MA: MIT Press.
- Recanati, F. 1994: "Contextualism and Anti-contextualism in the Philosophy of Language", in Tsohatzidis, S, ed. *Foundations of Speech Act Theory*, 156-66. London: Routledge.
- Recanati, F. 1995: 'The Alleged Priority of Literal Interpretation', *Cognitive Science*, 19, pp. 207-232.
- Recanati, F. 2003: *Literal Meaning*. Cambridge: Cambridge University Press.
- Sperber, D. and Wilson, D. 1986: *Relevance: Communication and Cognition* (2nd edition, 1995). Oxford: Blackwell.
- Sperber, D. and Wilson, D. 1998: 'The Mapping between the Mental and the Public Lexicon', in P. Carruthers & J. Boucher, (eds) *Language and Thought: Interdisciplinary Themes*, pp. 184-200.
- Travis, C. 2000: *Unshadowed Thought*. Cambridge, MA: Harvard University Press.
- Wettstein, H. 1979: "Indexical Reference and Propositional Content", *Philosophical Studies*, 36: 91-1000.
- Wilson, D. 2003: Relevance and Lexical Pragmatics. *Rivista di Linguistica/ Italian Journal of Linguistics*, 15.2: 273-291.
- Wilson, D. and Sperber, D. 2002: 'Truthfulness and Relevance', *Mind*, 111: 583-632.