Word Meaning, What is Said, and Explicature[*]

Robyn Carston
Linguistics, University College London, and
Centre for the Study of Mind in Nature, University of Oslo

1.   Background: what is said, what is meant and what is explicitly communicated

In his logic of conversation, Paul Grice made an important distinction between two kinds of act that speakers may perform when uttering a sentence and, correspondingly, two kinds or levels of speaker meaning: (i) saying, hence what is said, and (ii) implicating, hence what is implicated (implicature) (Grice 1975). The distinction between what is said and what is implicated can be seen as one instantiation of a distinction between the explicit content of an utterance and its implicit import. This distinction, albeit with numerous modifications and extensions, has proved indispensable to all subsequent attempts to account for speaker meaning, communicated content and utterance interpretation.

It is generally agreed, however, that in his few characterisations of 'what is said', Grice underestimated the extent to which context/pragmatics can contribute to a hearer's recovery of this level of content. He maintained that 'what someone has said [is] closely related to the conventional meaning of the words (the sentence) he has uttered' (1975: 44), and, beyond conventional linguistic meaning, he explicitly acknowledged only the need to identify the reference of pronouns, the time of utterance, and the intended sense of any ambiguous words or phrases (1975: 44). In the meantime, a plethora of linguistic phenomena has been put forward as requiring or, at least, allowing for further contextual/pragmatic effects at the level of explicit utterance content. These include the following: (i) cases of syntactically complete but semantically non-propositional sentences, e.g. 'Jane is ready', 'Bill is tall enough'; (ii) cases of (apparently) unarticulated constituents of content, e.g. the location constituent usually inferred for utterances of 'It is raining', the cause-consequence constituent often inferred for conjunctive utterances such as 'Jane did Bill's lectures and he recommended her for promotion'; (iii) the modulation of word meaning, e.g. colour terms as in 'red bird', 'red house', 'red pen', 'red crystal', 'red hair', and so on; (iv) more controversially, certain kinds of non-literal uses (metaphorical or metonymic) of words, e.g. 'John is a saint', 'Mary is a butterfly', 'The mushroom omelette wants her bill'.

It has been suggested that, if he had considered the matter in more detail, Grice would not have been averse to including at least some of these kinds of pragmatic contributions in what is required for a full identification of what the speaker has said (Wharton 2002, Stephen Neale (pc)). This may be so for semantically incomplete sentences as in (i) above: Grice might have agreed to add to what is required for a full identification of what

the speaker has said the answers to such questions as 'ready to do what?', 'tall enough for what purpose?'. However, several other aspects of his work strongly indicate that he would not have made any substantial move in this contextualist/pragmaticist direction: (a) he was well aware of many of the phenomena now treated as pragmatic contributions to explicitly communicated content, e.g. the various non-truth-functional enrichments of conjunctions and disjunctions, cases of scalar inference and the referential use of definite descriptions, and he explicitly analysed them as cases of conversational implicatures (Grice 1975, 1978); (b) for metaphor and other non-literal uses, he again took an implicature approach and, most tellingly, moved to talk of 'making as if to say' rather than allow anything other than literal word meanings to enter into 'what is said'; and (c) the whole motivation for his saying/implicating distinction was to keep meaning (semantics) distinct from use (pragmatics) and to ensure that judgements of truth/falsity were confined to the former (Grice 1967/89: 3-21; for discussion of this point, see Carston 2002: 101-105).

What Grice wanted from his notion of 'what is said' was that it serve both as the semantic content of the sentence uttered and as the proposition asserted (hence meant) by the speaker. However, as I've argued at length elsewhere, no single level of meaning can do double duty as both sentence semantics and speaker-meant primary meaning[1] (Carston 2004; Carston & Hall 2012). This is evident in the work of virtually all those who see themselves as, in some sense, following, developing or reacting to this area of Grice's thinking. His notion of 'what is said' has had to be reshaped and, in most instances, split into two distinct components, one semantic, the other pragmatic. For instance, Kent Bach has opted for a strict semantic notion of 'what is said', which is not speaker-meant, and has posited a distinct level of speaker-meant content, called 'implicirue', which is the result of pragmatic processes of completion and expansion of the minimal what is said (Bach 1994). Francois Recanati and other contextualists maintain a pragmatic notion of what is said, which is speaker meant, but accept that this may require much more elaboration and modulation of the linguistic meaning than Grice allowed (Recanati 2004), and they generally take it that sentence semantics is much more schematic than Grice envisaged.

Relevance theorists have argued that Grice's notion of 'what is said' (and the label itself) should be dropped as it is an artificial construct that doesn't pick out any psychologically real component or level of meaning (Carston 2002, 2004; Wilson & Sperber 2002). Like Recanati, they maintain that linguistically-given meaning (semantics) is schematic (non-propositional), functioning as evidence for (or constraints on) the propositional contents that a speaker has communicated, and that the primary communicated content which a hearer must reconstruct pragmatically from this linguistic basis is generally considerably richer than Grice's 'what is said'. Sperber & Wilson (1986/95: 182) introduced the term 'explicature' (as a counterpart to 'implicature') for this primary component of speaker

---

[1] The phrase 'primary meaning' is used in an attempt to find a neutral terminology, given the various arguments in the literature over the pros and cons of terms like 'explicature', 'implicirue', and an intuitive or pragmatically enriched 'what is said'. 'Primary meaning' is the content meant (or communicated) by a speaker which is developed on the basis of the linguistic meaning and it is distinct from 'implicature' (the speaker's 'secondary meaning').

meaning, which is a pragmatic development of the linguistically decoded content.[2]

An utterance of the sentence in (1a) below could have any of indefinitely many possible explicatures, depending not only on the intended referents of the pronouns, but also on the intended meaning of 'take' and the intended completion of 'enough'. And the same goes for (1b) with regard to what is meant by the phrasal verbs 'get through' and 'take out' and what the intended completion of 'get through' is on a specific occasion of use.

(1)   a.   She has taken enough from him.
          b.   When he gets through she'll take him out.

The content of explicatures (as derived by hearers) comes from two distinct sources, the linguistic expressions used and the context, and it is derived in two distinct ways depending on its source, by linguistic decoding or by pragmatic inference. It follows that explicatures can vary in their degree of explicitness depending on the relative contributions of decoding and context, so although an utterance of the sentence in (2) could have the same explicature as an utterance of the sentence in (1a), the contribution from linguistic encoding is greater and so it is more explicit:

(2)   She has taken enough help from Peter.

Although context (linguistic and extra-linguistic) plays a substantial role in the relevance-theoretic account (RT), the theory departs from existing 'contextualist' positions on the issue of how context enters into the derivation of explicatures. For instance, on Recanati's contextualist account, the derivation process is a matter of contextual best fit, achieved via direct modulation of linguistically decoded meaning by salient components of context (Recanati 2004). On the RT 'pragmaticist' account, on the other hand, the pragmatic enrichment and adjustment of linguistically provided content involves the exercise of full-blown pragmatic inference[3] and it is only (what are taken to be) intended contextual assumptions that enter into the interpretation process. However, for the purposes of this paper, it is what the two approaches share that matters most, that is, the view that there is typically a greater gap between linguistically given meaning and the speaker's primary meaning (explicature) than can be bridged by merely fixing the reference of indexicals and disambiguating homonymous lexical forms.

The issue I want to consider in this paper concerns the nature of word meaning, in particular the standing lexical meaning of words in what are known as the 'open classes',

---

[2] Setting aside terminological differences, the three positions just discussed (Bach's, Recanati's, RT's) are very similar, at least with regard to the levels of content envisaged: a minimal linguistic content (semantics), a primary speaker meaning, which is the result of pragmatic enrichments and adjustments of the linguistically provided content, and secondary speaker meaning(s), which are wholly pragmatic inferred (implicatures).

[3] By 'full-blown pragmatic inference', I mean the process of drawing sound conclusions from a set of premises, which may include premises concerning mental states of the speaker (her beliefs, intentions), rather than simply associative mappings and concept activation. For more on the distinction between contextualism and pragmaticism, see Carston (2009), and for a critique of some ways in which 'context' is invoked in accounts of pragmatic enrichment, see Bach (2005).

specifically, nouns, verbs, and adjectives. The general linguistic underdeterminacy claim concerns the relation between the explicature a speaker expresses by her use of a sentence on a particular occasion and the (standing, invariant) meaning of the sentence within the language. The move now is to consider the relation between the concept a speaker expresses by her use of a word on a particular occasion and the stable meaning that the word has as a lexical item in the language. The central question to be addressed is whether words encode concepts, hence the same kind of entity as they are used to express, or something more schematic, like a blueprint or template for concept construction. If it turns out that the latter is the case, then we would have a nice parallel between words and sentences, both of them underdetermining of what they are used by speakers to express. I don't have a definitive answer to offer to the difficult question of what context-independent word meaning consists of, but will provide some considerations that support the position that word meaning, like sentence meaning, is inevitably underdetermining of what it expresses and that it is different in kind (it is non-conceptual).

The second question is what view of standing word meaning is to be expected or preferred within the relevance-theoretic account of utterance interpretation. According to this account, word meanings are pragmatically modulated in context, resulting in ad hoc concepts which are components of the conceptual representation recovered by the hearer as the proposition explicitly communicated by the speaker (the explicature). I will argue that, while minimalist propositional views of semantics (including Grice's semantic conception of 'what is said') seem to require that word meanings are fully conceptual, the pragmaticist position of RT does not require this and seems in many respects better served by a more schematic view of encoded word meanings.

The rest of the paper is structured as follows. In section 2, I present the account of word meaning in current relevance theory, both its view of encoded (or 'standing') word meaning and its account of 'lexical pragmatics', according to which a word often contributes not its encoded meaning but an ad hoc concept to the explicature. In sections 3 and 4, which are the core of the paper, I set out some arguments against construing encoded word meanings as concepts and outline some psycholinguistic evidence which supports an underspecification view. Then, in section 5, after a brief look at the account of word meaning entailed by semantic minimalist accounts, I return to the relevance-theoretic account of word meaning and suggest that it should embrace a non-conceptual position. Finally, in the conclusion, section 6, I relate the position on word meaning taken here to a broader stance concerning the inevitability of linguistic underdetermination of speaker meaning and the impossibility of 'saying' what one means.

2.   Relevance theory: word type meaning and occasion-specific word meaning

A distinction between two kinds of word type meaning (or lexical semantic encoding[4]) is made in relevance theory: conceptual encoding and procedural encoding (Blakemore

---

[4] I am aware that for many people 'lexical semantics' is associated with a concern to capture sense relations between words (synonymy, hyponymy, antonymy, etc.) and so usually involves detailed (decompositional)

1987, 2002; Wilson & Sperber 1993; Wilson 2011). Diane Blakemore first introduced the idea of procedural meaning in the context of a class of terms she called discourse connectives, e.g. 'but', 'moreover', 'after all', 'so', 'you see', which arguably do not contribute to the truth-conditional content of utterances in which they occur. She claimed that the lexically encoded meaning of such words is not a concept (a component of content) but rather a constraint on the kind of pragmatic inferential process the hearer is to undertake in deriving the speaker's meaning. Since that early work of hers, the notion of procedural meaning has been developed and the range of linguistic elements it is claimed to apply to has been extended (e.g. pronouns, particles, tense, aspect and mood markers, and interjections have all been claimed to encode procedures). My focus in this paper is on those words that are taken to encode concepts, that is, open-class words like nouns, verbs and adjectives, so I am mentioning procedural meaning at this stage in order to set it aside, although the notion will reappear in section 5 in the context of a general non-conceptual view of word meaning.

According to the RT view, the concepts encoded by words (lexical concepts) are unstructured, that is, they do not decompose into components that comprise definitions, or stereotypical features, or prototypes, and their content is not dependent on their embedding in any theory or system of beliefs. A semantic representation of a natural language sentence is a syntactically structured string of atomic (or 'primitive') concepts (together with some elements of procedural meaning as mentioned above). So, as regards the meaning/semantics of words, the mental lexicon is very minimalist, simply providing mappings from lexical forms to simple concepts, e.g. 'red' maps to RED; 'open' maps to OPEN; 'cat' maps to CAT, a position that is much akin to Jerry Fodor's 'disquotational' lexicon (Fodor 1998).[5] These concepts have a denotation, that is, they refer to entities, properties and activities in the mind-external world (e.g. cats, the property of redness, the action of opening) and, as mental representational entities, they play (at least) two roles in our cognitive life. They are the basic constituents of our language of thought or Mentalese, the system of mental representation in which we think and in which the explicatures and implicatures of our utterances are couched. And each such primitive concept (RED, CAT, OPEN, etc.) is a storage point for encyclopaedic information,[6] including general knowledge and individual beliefs about the things it denotes, cultural knowledge, including stereotypes, which the individual may or may not endorse, imagistic representations, and perhaps also episodic memories. (Much more needs to be said about how this information is organised and tagged so that general knowledge, stereotypes, individual memories, etc. are kept distinct from each other.) Encyclopaedic information, in particular general and cultural knowledge, is a major source of contextual assumptions, which play a central role in the

---

analyses of word meanings. As will be clear, this is not the sort of semantics for words that is advocated here.

[5] However, RT does not fully embrace Fodor's conceptual atomism because it maintains a partially inferential role semantics for concepts in the form of meaning postulates such as 'α RED → α COLOURED', 'ϕ CAT ψ → ϕ ANIMAL ψ' (see Sperber & Wilson 1986/95 and, for a defence of meaning postulates against Fodor's objections, see Horsey 2006).

[6] There are various metaphors for what this amounts to, including the idea that concepts are addresses in memory (Sperber & Wilson 1986/95: 86) or labels for files of information.

relevance-theoretic account of utterance interpretation and specifically in the process of lexical concept adjustment, which I will go on to outline now.

According to recent work on 'lexical pragmatics' within relevance theory, a word's meaning may be narrowed or broadened on occasions of use, that is, the concept a speaker communicates may be more specific or more general than the concept the word encodes (Wilson & Carston 2007). These pragmatically inferred concepts are known as 'ad hoc concepts' in the theory - they are distinct from lexically encoded concepts and are occasion-specific. An ad hoc concept derived from an encoded concept is marked with an asterisk, e.g. RED* is a concept pragmatically inferred from the lexical concept RED on an occasion of use, OPEN* from OPEN, etc. Lexical narrowing and broadening are to be understood in denotational terms, that is, the denotation of the communicated concept may be a subset or a superset of the set that comprises the denotation of the encoded concept. A combination of the two is also possible so that the denotation of the resulting concept may merely overlap with that of the lexical concept. These various outcomes are the result of a single process of lexical adjustment that takes place as part of the overall inferential pragmatic process of utterance interpretation. Consider the following examples, focussing on the highlighted word:

(3) a. There's *milk* in the fridge.
     b. I've given up *teaching*.
     c.   Chris *ran* for Britain.

In appropriate conversational situations, each of these could be understood as communicating a concept which is more specific than the concept it encodes: in (3a), the milk referred to might have to be of a certain quantity and quality making it suitable for drinking or adding to coffee (a few stale drips on a shelf would fall outside the denotation of MILK*, the concept communicated); in (3b), it might be evident that the speaker means she has given up classroom teaching (she might still be actively involved in teaching her grandchildren how to make pancakes or her neighbour how to prune a magnolia); in (3c), the communicated concept RUN* might denote only those instances of running that meet the standards of athletes competing in international races.

The sentences in (4) are all instances of literal uses of the transitive verb 'pass' (without any of the prepositional or adverbial additions that might make for a different verb, e.g. 'pass on', 'pass out', 'pass away', pass into', 'pass off', 'pass over', 'pass down', etc.) and yet the concept expressed and understood in each case is likely to be distinct from the others:

(4) a. Jack passed a rope around the tree.
     b. Jill passed two cars.
     c. Jack passed Jill his phone number.
     d. The dog passed wind.
     e. She passed her exams.
     f. Eventually his depression passed.
     g. They passed the time happily.

If each of these communicates a more specific concept (PASS*, PASS**, PASS***, etc.) than the concept PASS encoded by the lexical form 'pass', the question that arises is just what this more general concept is and whether it is ever communicated. Perhaps what we have here is a case of polysemy, that is, of one word with a number of related senses, or of some stage prior to semantic polysemy, as it seems unlikely that each of these concepts is a fully conventionalised sense which has been incorporated into the lexicon. These and other issues raised by this sort of case are considered in more detail in section 3.

The following are cases where pragmatic lexical adjustment of the italicised word would probably have the opposite result, that is, an ad hoc concept with a broader denotation than the encoded concept:

(5)   a.   The children quickly formed a *circle*.
         b.   This steak is *raw*.
         c.   She *raced* down the street.
         d.   That lake is an *ocean*.

If the word 'circle' encodes the concept paraphraseable as 'closed flat curve every point of which is equidistant from the centre' (which might be disputed, of course), then the concept CIRCLE* communicated by an utterance of (5a) is likely to encompass a much broader range of shapes. Many precise terms (e.g. words for numbers, geometric shapes, distances, and times) are often used in this sort of approximate way. Other words may be used hyperbolically in order to have a particular forceful effect, as is very likely the case with 'raw', 'race' and 'ocean' in (5b)-(5d), and the concepts thereby communicated will clearly be broader than the concepts they encode.

According to orthodox relevance theory, metaphorical uses of words, as in (6), are also cases of concept broadening, so, for instance, the denotation of the concept CIRCULAR* communicated by an utterance of (6a) includes not only instances of a certain kind of physical shape but also certain patterns of thoughts and the concept OCEAN* communicated by (6d) would encompass not only large bodies of salt-water but also human feelings of a particular strength and constancy:

(6)   a.   Your argument is *circular*.
         b.   His grief was still *raw*.
         c.   His thoughts were *racing*.
         d.   Her love for him was an *ocean*.

It might well be that this is only part of the story for metaphorical uses, which many theorists think of as involving a domain shift (e.g. from the physical to the psychological) and so as involving a denotational narrowing as well as a broadening (or some other process altogether). I won't pursue this issue here,[7] as the relevant point for current purposes is the

---

[7] For the orthodox relevance-theoretic view on metaphor, see Wilson & Carston (2006) and Sperber & Wilson (2008); for a modification of this view, see Carston (2010) and Carston & Wearing (2011).

general one: pragmatic lexical adjustments of one sort or another occur frequently within the process of utterance understanding and the concept encoded by a lexical form is merely a clue to (or a constraint on) the intended interpretation.

As my main interest in this paper is the nature of encoded word meanings (of open class words), that is, the lexical input to the pragmatic process of grasping the concepts communicated by speakers' uses of words, I will not go into the detail of the relevance-theoretic comprehension process but give here just a very brief sketch of its main components (for more thorough accounts, see Wilson & Sperber 2002, 2004, Wilson & Carston 2007).

When a lexical concept is accessed via the usual linguistic decoding process, the encyclopaedic information associated with it is activated. Some elements of it are more highly activated than others (since there are multiple sources of spreading activation, including other concepts encoded in the utterance and conceptual representations derived from the wider discourse or situation of utterance). The most highly activated items of conceptually represented information are accessed and deployed as contextual assumptions in deriving contextual implications, which form an initial interpretive hypothesis about the utterance. Then, via a mechanism of mutual parallel adjustment of explicit utterance content (explicature), contextual assumptions and contextual implications, lexical concepts in the decoded semantic representation are adjusted by backwards inference, so that only implications that are ultimately grounded in the explicature are confirmed. The overall interpretation is accepted provided that it meets the addressee's expectation of relevance. Consider example (3a) '*There's milk in the fridge*' uttered in a situation in which the speaker and hearer are about to have a cup of coffee. The utterance would have the contextual implication that there is milk available to add to their coffee, which would be supported by a contextual assumption (a highly accessible item of encyclopaedic information) that milk comes in cartons and bottles for a variety of culinary uses including adding to tea and coffee. By backwards inference, the concept MILK* is formed, paraphraseable as 'milk of sufficient freshness and quantity to be suitable to add to coffee' and this concept is a component of the explicature. In another utterance situation, for example, one in which the hearer is supposed to have meticulously cleaned the fridge, different contextual assumptions and implications (specifically, that there is still some mess to be cleaned up) could result in a distinct ad hoc concept MILK**, paraphraseable as 'drops of stale milk'. And there are numerous other possibilities.

In the next two sections, I focus on the theory-general question of what sort of thing stable (lexicalised) word meanings are, with a view to making a preliminary case for their being non-conceptual (non-semantic). I then return to RT in section 5 and argue that this non-conceptual view of word meaning is compatible with the theory, perhaps even preferable to the conceptual view.

3. Words behave as if their meanings are conceptually underspecified

The view that stable (encoded or standing) word meanings are concepts, i.e. mental representations with a denotation, is fairly widespread: it is held by philosophers including

Fodor (1998), Fodor and Lepore (1998), Borg (2012) and probably others who fall under the label of 'semantic minimalists', by psychologists including Bloom (2000), Murphy (2002) and Hampton (2011), and by relevance-theoretic pragmaticists, as discussed in the previous section. Others, mostly linguists, have argued for some kind of semantically underspecified meaning.[8] These include Charles Ruhl (1989), who advocates the 'monosemic bias' in the face of words which seem to have many (potentially indefinitely many) related senses, a stance that leads inevitably to a take on standing word meanings as underspecified (that is, as not constituted by any one of the specific senses that is expressed in use but of something more abstract that is shared by all of them).[9] In recent work within the Dynamic Syntax framework, Ruth Kempson and her colleagues (2011) have characterised word type meanings as 'lexical actions' or procedures, which together with the instructions provided by the syntax of a language constitute a set of mechanisms enabling hearers to construct representations of content. That construction process is heavily reliant on processes of pragmatic inference, such as those discussed in section 2, with the lexical procedures functioning as constraints on inference.

The computational linguist, Peter Bosch, maintains that the lexical type meaning of many words is underspecified and must be developed at a conceptual/pragmatic level in order for its expressed meaning to be recovered by a hearer or processor. The noun 'novel', for example, discussed at length in Bosch (2007), can have the following senses: a complex of ideas/thoughts (when the author is working on it), a text (when it is completed), a publication (e.g. when we talk of an author's most recent novel), a physical object (e.g. when we talk of a suitcase full of novels), and certain combinations, e.g. '*Peter is reading the novel he found at the bus-stop*' (text and physical object). No particular one of these senses is obviously the encoded meaning or is sufficiently all-encompassing to provide the basis for pragmatically inferring the other senses. As Bosch says: 'If we want to maintain just one lexical entry for *novel* it must remain underspecified in many respects …' (2007: 68). He discusses a number of other words of various lexical categories (e.g. 'work', 'rain', 'run', 'cut', 'open', 'fast') whose susceptibility to being used to express different concepts, which he calls 'contextual concepts', points to the same conclusion: 'the lexical semantics should be left underspecified in these cases' (2007: 58). However, Bosch talks throughout the paper of the stable standing meaning of a word as a lexical *concept* and proposes to follow the 'disquotational' view of lexical semantics advocated by Fodor & Lepore (1998). In his discussion of the word 'novel', for instance, he says:

> '… the various senses each represent a different *view* of our underspecified concept NOVEL. …. The concept NOVEL may be subsumed by concepts like PHYSICAL_OBJECT, TEXT, PUBLICATION, etc. In each such view the resulting

---

[8] A few philosophers also support some kind of underspecification position on word meaning. Charles Travis, famous for demonstrating the different ways in which an unambiguous word (e.g. 'green', 'cover', 'sailor') may be understood on different occasions of use, talks of words as providing constraints on those different understandings (Travis 2000: 221-227). Tim Pritchard is actively engaged in trying to provide an account of core stable meanings for words that differ in kind from the senses (concepts) they are used to express (Pritchard forthcoming).

[9] Ruhl's study of the verb 'bear', in which he provides a wealth of attested examples, brings out the range of fine-grained differences in sense that the verb can be and has been used to express (Ruhl 1989, chapter 2).

contextual concepts for *novel* (in all cases sub-concepts of the lexical concept NOVEL) inherit different attributes from the different superconcepts under which they are being subsumed.' (Bosch 2007: 68).

The idea that underspecified word meanings are lexical concepts raises a question for me. While lexical concepts of the Fodorian sort are minimalist in many respects (no internal structure, no content-constitutive inferential connections or knowledge structures), they are full-fledged concepts nonetheless: they are components of thoughts ('words' of Mentalese) and have a referential semantics. So the question is: what is the lexical concept that comprises the stable meaning of a word whose use gives rise to a range of distinct contextual concepts (or 'ad hoc concepts', in relevance-theoretic terms)? What concept is the lexical concept NOVEL as distinct from the various contextual concepts NOVEL*, NOVEL**, etc.? These latter concepts are clearly constituents of thoughts that we have and are able to communicate, but what about the lexical concept, a concept which is somehow neutral as regards novels as physical objects, narratives, texts, publications - does it occur as a constituent of a thought, and, if it does, what is that thought about, that is, what is the denotation of NOVEL? The same sort of issue arises for every word that can be used to express myriad different concepts on different occasions: for instance, what is the (alleged) general lexical concept OPEN and what is the content of a thought in which such a general concept features, as opposed to any of the more specific concepts that we grasp in understanding *open one's mouth*, *open the window, open a letter, open a discussion*, etc.? Finding no ready answer, the question then becomes: is there any definite thought at all that engages the general concept or must any thought about opening contain one of the more specific concepts? (See Carston 2012 for a little more detail on this point). The difficulty in getting any purchase on these questions about underspecified lexical concepts is a first indication that we might be well advised to move to a different, non-conceptual (non-semantic) view on the nature of standing word meanings. As Bosch puts it in a later paper: '… lexical semantics is much less specified than is often assumed and only contains *structural constraints* over the kind of conceptual entities that can be denoted by the lexical item, but *does not contain the conceptual content*.' (Bosch 2009: 99; my emphasis)

Sperber and Wilson (1998) maintain that some words do not encode a full-fledged concept but what they call a 'pro-concept', that is, a meaning or character that functions as a constraint on the kind of concept that it can be used to express. What that concept is on any occasion of use has to be pragmatically inferred using relevant contextual information. They suggest 'my', 'have', 'near', and 'long' as likely instances of words that encode pro-concepts, and say that 'while each of these examples may be contentious, the existence of the general category should not be' (Sperber & Wilson 1998: 185). Most interestingly, they go on to say: '… quite commonly, all words behave *as if* they encoded pro-concepts: that is, whether or not a word encodes a full concept, the concept it is used to convey in a given utterance has to be contextually worked out' and they discuss the verb 'open' as a case in point: 'A verb like 'open' acts as a pointer to indefinitely many notions or concepts ...' (1998: 197). However, they also maintain that the pragmatic process of inferring the intended concept in context is optional because 'it may so happen that the intended concept is the very one encoded by the word,' (1998: 197). So Sperber and Wilson presup-

pose that many words encode concepts[10] as their standing meaning and may, on occasion, contribute the concept they encode directly, without any pragmatic adjustment, to explicatures. This view is carried over into more recent RT work on lexical pragmatics too, even though this work has led to the claim that '… lexical narrowing and broadening (or a combination of the two) are the outcomes of a single interpretive process which fine-tunes the interpretation of *almost every word*.' (Wilson and Carston 2007: 231, my emphasis). If words quite generally behave as if they don't encode concepts, we might wonder why we should assume they do encode concepts, as opposed to something more schematic, which merely constrains or guides the pragmatic process of recovering the concept a speaker has expressed.[11]

Any move away from a fully conceptual standing meaning for words will, of course, have a range of consequences and raise new questions. One of the most pressing of these concerns the pragmatic inferential process of constructing ad hoc (contextual) concepts, a process which would seem to be obligatory if the lexical input is underspecified. The question is how, given a starting point that is non-conceptual, hence non-semantic, does the pragmatic process result in a semantic (truth-conditional) output.[12] Confronting and answering this question requires more space than I have here and I defer it to a later paper.

Another possible concern is whether this move might fall foul of the fundamental principle that the meaning of the whole (the sentence or thought) should be a determinate function of the meanings of the basic constituents (word meanings) and their mode of combination, that is, the principle of semantic compositionality. Fodor (1998) maintains that this is a non-negotiable requirement on an account of concepts and he has used it as the key argument in support of his particular view of lexical concepts as referential atoms and against other views of concept content (as prototypes, inferential roles, theories or bodies of knowledge/beliefs). The principle is often presented as applying equally to concepts/thoughts and word meanings/sentences, but it is a principle of *truth-conditional* semantic compositionality, so if sentence meaning is not truth-conditional this particular principle of compositionality does not, after all, apply to language. That sentence meaning quite generally underdetermines truth-conditional content is a central tenet of contextualism/pragmaticism, as discussed in section 1. In other words, natural language sentences are simply not compositional in the required sense, as Fodor himself occasionally acknowledges: '… a perfectly unelliptical, unmetaphorical, undeictic sentence that is being used to express exactly the thought that it is conventionally used to express, often

---

[10] This is made especially clear in Wilson (2003: 274; my emphasis): 'I will adopt a simple model of linguistic semantics that treats *words as encoding mentally-represented concepts*, elements of a conceptual representation system or 'language of thought', which constitute their linguistic meanings and determine what might be called their *linguistically-specified denotations*.

[11] Note that a non-conceptual account of the open-class vocabulary will not be one that adopts wholesale indexicality (or 'pro-concepts', a notion that seems to entail an indexical component). This point is elaborated on a little in section 5 of the paper.

[12] The lexical concept view of word meaning advocated by Sperber & Wilson, and apparently also by Bosch (2007), does not raise this question: the pragmatic adjustment process takes a semantic entity as input (a concept) and outputs another semantic entity (a distinct concept). On the face of it, this is a definite advantage of the conceptual account.

doesn't express the thought that it would if the sentence were compositional ….' (Fodor 2001: 12) and he goes on to talk of the typical case of this situation being one in which a sentence 'vastly underdetermines the right thought', concluding that 'The evidence strongly suggests that language is not compositional' (2001: 4). It follows, then, that the primitives of natural languages (words or morphemes) are not bound by the compositionality principle that constrains an account of primitive concepts (the 'words' or basic components of the language of thought). The standing context-invariant meanings of words need not contribute thought components (concepts) unmediated by pragmatic processes and there is, therefore, no obstacle here to their being some kind of non-conceptual entity, which merely places constraints on the concepts that a word can be used to express.

There is a further point worth noting, though, concerning the motivation for the compositionality requirement. Language and thought both have the properties of 'productivity' (unboundedness of the number of distinct representations they can generate) and 'systematicity' (sameness or overlap of constituent parts of sets of representations, such as, for example, the set of thoughts/sentences 'Dogs chase cats', 'Cats chase dogs', 'Do dogs chase cats?'). It is widely (although not universally) agreed that it is the compositionality of mental and linguistic representations that underlies these two properties (for extensive discussion of this point, see Fodor 1998: 94-100; Fodor & Lepore 2002). However, the (undoubted) productivity and systematicity of a natural language are accommodated by a fairly trivial notion of the compositionality of sentence meanings, one which follows directly from its recursive syntax coupled with stable meanings for its primitives (words or morphemes). So long as we maintain that sentence meanings are not truth-conditional entities (thoughts/propositions) and that pragmatics plays an essential role in deriving the propositions expressed by speakers, the constraints that productivity and systematicity (specifically the latter) place on our account of word meaning are meagre: that those word meanings are constant wherever they occur in a linguistic representation. As far as I can see, there is no reason to suppose that a non-conceptual account of word meaning would be any less able than a conceptual account to meet this requirement.

In this section, I have presented some arguments in favour of an account of word meanings as underspecified and non-conceptual. None of them provides an overwhelming case and the idea needs a great deal more thinking through. However, there is some concrete empirical evidence that appears to be pointing in the same direction. I turn to that evidence in the next section.

4.   Polysemy and evidence for underspecified word meanings

The work I am going to review here concerns the processing of words that have the property of being 'polysemous', that is, broadly speaking, words that have multiple related senses (where that relatedness is apparent to users of the word). There seem to be at least two ways of viewing the phenomenon of 'polysemy', a pragmatic way and a semantic way. The various examples discussed in preceding sections, including 'red', 'milk', 'teaching', 'circle', 'raw', and the verbs 'pass', 'open', 'bear', have all been shown to express different related meanings on different occasions of use. With regard to the process-

es of utterance comprehension, this is arguably an 'output' phenomenon, the result of pragmatic processes of interpretation (hence *pragmatic* polysemy). There is another, probably more common, way of thinking about polysemy, which is as an 'input' phenomenon and so as applicable only to those instances where the distinct but related senses are fully conventionalised (hence *semantic* polysemy). Examples often cited are the animal/meat alternations for words like 'lamb', 'chicken' and 'rabbit', cases like 'book', 'novel', 'film', 'dvd' which are regularly understood as either the physical object or the representational content, and numerous others where a metaphorical meaning has become very common and alternates with a literal meaning, e.g. adjectives like 'cold', 'warm', 'hard', 'soft', 'rigid', 'flexible', 'bright', 'dull', and so on.

Pragmatic polysemy is surely the forerunner and source of semantic polysemy; while many cases of pragmatic polysemy will not become cases of semantic polysemy (many ad hoc concepts are one-off or at least too infrequent to become established senses), all cases of semantic polysemy must have arisen pragmatically (even if in some instances we are unsure which sense came first). A different way of putting what is essentially the same point is the following from Fodor & Lepore (2002: 117):

> 'Suppose it's right that "lamb" is polysemous between the animal and the meat. Surely that's because lamb-the-meat comes from lamb-the-animal. Surely there just *couldn't be* a word that is polysemous between *lamb-the-animal* and (say) *beef-the-meat*? Or between *lamb-the-animal* and *succotash-the-mixed-vegetable*? … Opportunities for polysemy arise from how things are in the world (or, anyhow, from how we take them to be).'

Speakers are able to exploit these opportunities because they can rely on the propensity of words, when mentally accessed, to activate associated bodies of world knowledge/assumptions, thereby enabling hearers to pragmatically construct relevant concepts that are distinct from, but related to, the concept(s) directly associated with the word.[13]

Let me turn now to the experimental work on polysemy by the psychologist Steven Frisson and his colleagues, which is of considerable interest for the hypothesis being considered here: that standing word meanings are non-conceptual (non-denotational). The point of departure for these polysemy studies was the existing abundant and robust experimental results on the processing of homonymous lexical forms, e.g. 'bank' (with the two meanings: (i) a financial institution, (ii) a side of a river), 'coach' (with the two meanings: (i) an instructor, (ii) a large motor vehicle), that is, lexical forms which happen to map to two (or more) senses which are unrelated (at least as far as the language user is concerned[14]). Summarising these results, when a hearer encounters such a form both mean-

---

[13] See also Bosch (2009), who maintains, similarly, that 'polysemy has very little to do with lexical semantics and compositionality. Polysemy is neither in the lexicon, nor does it arise from an interaction of different words in a sentence. Instead, … polysemy effects come about when lexical information interacts with the denotation that an occurrence of a particular expression has in a particular context. Such interactions yield context-specific denotations for expressions, and it is these denotations that are the input for compositional processes, which eventually will give us something like the "meaning" of an utterance.' (p.101).

[14] It could be that the meanings of some homonymous lexical forms are, in fact, etymologically related, something that would be noted in a well-researched dictionary, but this is irrelevant to their psychological

ings are briefly activated initially (but ranked according to frequency) and a decision is reached rapidly in favour of one or the other (on the basis of preceding context or, if the context doesn't help, on the basis of frequency). In general, when the subordinate (less frequent) meaning is intended, more processing is required; in particular, when only *subsequent* context indicates that the subordinate meaning is intended, a great deal of extra processing occurs, indicating that the wrong meaning had been selected initially. In short, language users make an early semantic commitment in that they immediately select one or other of the meanings. (For more detail, see Frisson & Pickering 2001; Pickering & Frisson 2001; Frisson 2009.)

In a series of experiments[15], Frisson and colleagues have found evidence that the senses associated with polysemous words are processed differently from those associated with ambiguous lexical forms (homonyms)[16] and the evidence points to what they call the 'Underspecification hypothesis', which is why it is of interest here. The cases of polysemy they have tested are those that would generally be judged as conventional or well-established, so nearer to the semantic end of the spectrum than to the one-off ad hoc concept end. These include nouns with a concrete and a more abstract sense like 'school' (the building and the institution), and those with a figurative sense, whether metonymic as in (7a)-(7b) or metaphorical as in (7c)-(7d):

(7)   a.   The blasphemous woman will have to answer to the *convent*.
        b.   A lot of Americans protested during *Vietnam*.
        c.   He got on his bike and *flew* down the road.
        d.   The speaker *disarmed* his critics with his wit and charm.

What they found, in summary, was that polysemous words do not show an effect of the relative frequencies of the two senses, and no extra processing effort is found when it is the subordinate sense that is the relevant one (even when it is only subsequent context that makes this clear). In contrast with homonyms, polysemous words are not processed in ranked parallel fashion, that is, the different senses are not all activated at the same time with the strength of activation dependent on sense frequency followed by rapid selection on the basis of frequency or context. Rather, no immediate semantic commitment is made, that is, language users do not immediately select one or other of the senses (Frazier & Rayner 1990, Frisson 2009). For instance, when processing the word 'flew' in (7b), people

---

status. If the individual language user stores the senses as two distinct, semantically unrelated, words then it is a case of homonymy (rather than polysemy).

[15] There is no space here to go into the details of the experimental designs and techniques used. Suffice it to say that the main methodology was eye-tracking, that is, the recording every few milliseconds of participants' eye movements as they read texts, including the length of their eye-fixations on words and their regressive eye movements to earlier sections of text. This technique enables the detection of those portions in a text that require more cognitive effort to understand than others.

[16] This might seem unsurprising to those of us who see the two phenomena as radically different, the one a matter of accidental coincidence of form (phonological and/or graphological), the other transparently pragmatically motivated. However, it has been claimed by other psycholinguists that polysemous words are in fact processed in the same way as homonyms and that, therefore, the different senses of a polysemous word are represented separately in the lexicon just like the distinct meanings of a homonymous form. For discussion, see Frisson (2009: 118-120).

do not immediately access the two senses, the literal one which is the dominant (more frequent) sense and the non-literal subordinate one, and rapidly select the latter on the basis of the context (riding a bicycle). Rather, the experimental evidence indicates that a single general meaning is accessed and this subsequently evolves into one or other of the more specific senses. On the basis of these findings (across a fair number of experiments testing different categories of words), Frisson and colleagues propose their 'Underspecification hypothesis':

> 'Rather than activating one or more specific senses, readers initially activate a single, semantically underspecified, meaning. This abstract meaning is the same for all established senses of a word, that is, the same underspecified meaning encompasses all semantically related interpretations of a word that are known to a reader.' (Frisson 2009: 116)

Once this underspecified meaning has been accessed, it can be followed by what they term a 'homing-in' stage in which context is used to arrive at one or other of the specific senses (see Frisson & Pickering 2001: 149; Frisson 2009: 117).

So the evidence from tracking people's on-line language processing provides support for the position that polysemous words have a single underspecified abstract meaning representation distinct from any of the senses (or concepts) that comprise the polysemy. The question might then seem to be: what about words that are not polysemous? After all, interesting though these underspecification findings are, if they apply to only a subset of the vocabulary, their relevance to the central issue of this paper is limited. My claim in response to this would be that *all* open-class words are either polysemous or at least potentially polysemous – that is the lesson to take from the pragmatic polysemies of section 2. In fact, even some closed-class words are polysemous (see Tyler & Evans (2001) on prepositions) and so are some derivational affixes, e.g. '-er', '-age', '-ery' (see Lieber (2004: 179), who talks of these as 'semantically underdetermined' affixes, whose semantic contribution may be realised in a number of ways depending on context). In short, polysemy is rampant and ubiquitous throughout the (non-syntactic) meaning-bearing elements of language. Evidence for an underspecified meaning for polysemous words is, then, evidence for an underspecification account of the meaning of (open-class) words quite generally, so the findings reported here are entirely germane to the hypothesis that word type meanings are not fully conceptual.

It is worth being clear about terminology at this point: I am using 'meaning' for the underspecified (non-conceptual, non-semantic) standing meaning of a word and 'sense' for the fully conceptual semantic entities that words are used to express. Thus, I would suggest the following rewording of the statements above from Frisson and Pickering about the making of 'semantic commitments' by hearers when processing words: 'An immediate *meaning* commitment is always made (for both homonymous and non-homonymous lexical forms), but an immediate *semantic* commitment is not made.' This is best illustrated with a lexical form which is both homonymous and polysemous: the lexical form 'bank' maps onto two distinct 'meanings' (unrelated and underspecified), that is, there are *two distinct words* with the same phonological/orthographic form. Both of these words are polysemous (or at least potentially polysemous), e.g. the word 'bank-1' has an underspeci-

fied meaning to do with financial matters and is associated with various different fully-specified senses including the bank as a financial institution, as a physical building, as a company, as a group of managers and other personnel (e.g. 'I talked to my bank this morning').

Finally, let's look briefly at the process that Frisson and Pickering talk of in terms of 'homing-in' on a specific sense.[17] They maintain that context is not to be seen as a 'judge' which selects among a number of senses/concepts but rather as an aid in developing a more specific interpretation of a word. This sounds entirely consistent with the view that I would advocate, as presented in section 2, that the process of arriving at the relevant (intended) sense of a word is pragmatic. However, the examples discussed in the work on lexical pragmatics involve the construction of ad hoc (occasion-specific) senses for which an inferential process is necessary, while many of the examples that Frisson and Pickering use are quite conventionalised, i.e. cases of 'semantic' polysemy (e.g. the two senses of 'school' or 'convent' as buildings and as institutions; the literal and the metaphorical senses of 'fly' and 'disarm'). It seems unlikely that homing in on these senses requires the kind of fully inferential pragmatic process involved in deriving ad hoc senses/concepts, but rather something more like a short-circuited inference or pragmatic routine (see Vega Moreno (2007) on pragmatic routines from a relevance-theoretic perspective). Nevertheless, the point remains that while the underspecified meaning is directly accessed or selected, the sense expressed, even when quite conventionalised, is a matter of contextual (i.e. pragmatic) specification of some sort. This is further supported by a discussion in Frisson and Pickering on the question of what is represented in the mental lexicon, whether it is both the single underspecified meaning *and* the various fully specified senses (which they call the 'radical polysemy' view) or just the underspecified meaning (which they call the 'radical monosemy' view). Their view is that: 'If we are correct that homing in on a specific interpretation involves using contextual information to refine an underspecified meaning rather than using context to guide the selection of an established sense, this would be more compatible with a monosemous account' (Frisson & Pickering 2001: 166-67). Prima facie, this has the air of paradox: polysemy is monosemy! But, of course, while 'polysemy' involves a plurality of *senses/concepts*, 'monosemy' applies to singularity of *meaning*, not of sense.

There is doubtless a great deal more to be said about polysemy, about its subtypes and, in particular, about how to accommodate the continuum of degrees of frequency/conventionalisation of the senses that cluster round a word. My primary interest here, though, has been to show that there is a body of empirical work on the processing of open-class words that seems to support the idea that standing word meanings are semantically underspecified, hence not fully conceptual.

---

[17] Frisson and Pickering do not think that 'homing-in' on a specific sense always takes place: in some situations, the more schematic non-specific meaning representation may be good enough. They say that whether a specific sense is accessed or not probably depends on several factors, including how important the word is in the sentence (e.g., whether it has sentence focus), how much contextual evidence there is for a specific interpretation, and what the requirements of the task are (e.g. whether there is time pressure or whether a full understanding of every single word is essential). See Frisson (2009: 117).

5.   Relevance theory and non-conceptual word meanings

Before revisiting word meaning within relevance theory, it is instructive to consider, if only very briefly, what a minimalist semantics position entails about word meaning. In clear distinction from the position of relevance theory (and other varieties of pragmaticism or contextualism), semantic minimalism maintains (i) that natural language sentences have a truth-conditional semantics, that is, well-formed declarative sentences semantically express propositions, and (ii) that a hearer's grasp of an utterance's truth conditions is achieved with minimal, if any, reliance on context, whether linguistic or extra-linguistic. Minimalists do recognise, as they must, that languages contain ambiguous forms and indexicals, which require that hearers look to context but, apart from these few cases, they take it that words are context-insensitive and their standing meaning makes a direct contribution to truth-conditional content. The inevitable upshot of this view is that the meanings of words must be concepts, that is, fully semantic entities which, when combined in the right way, are capable of yielding truth-evaluable content. As the staunch semantic minimalist Emma Borg puts it, 'the contribution words make to sentences [is viewed] *not* as open-ended, web-like things which stand in need of contextual precisification prior to fixing their input to larger linguistic units, but *rather* as discrete, probably atomistic, blob-like things.' (Borg 2012: xvii) She fleshes this out in terms of what she calls 'referential lexical axioms', that is, axioms that pair a word with an element of the non-linguistic world, e.g. the lexical axiom '*red* refers to the property of being red', adding that a full account would say that 'word meanings are concepts and it is these concepts which stand in relation to objects in the world/have their content given by things in the world.' (Borg 2012:144)[18]

Borg's minimalism is especially rigorous in that she insists that even demonstratives (e.g. 'this', 'that', 'there', 'then') contribute a context-insensitive semantic content (a singular concept) to the sentences containing them. Other semantic minimalists allow a role for context, even for pragmatic maxims and inference, in determining the semantic content of indexicals and demonstratives.[19] Nevertheless, the essence of all variants of minimalism is the same two properties: a propositional semantics for natural language sentences and very little context-sensitivity in the linguistic system. Open-class words, in particular, have a context-invariant semantics which entails that their standing meaning is conceptual, that is, they each contribute a determinate denotation to the truth-conditional content of every sentence in which they are a constituent. As discussed in section 1, I think that Grice was more minimalist than contextualist. He differs from contemporary minimalists in that he equated sentence semantics with what a speaker says, but, like them, he took sentence semantics to be propositional and, apart from disambiguation and fixing

---

[18] Borg (2012) tends to gloss over the role of concepts in the picture and to talk of words as relating directly to features of an external world because one of her primary concerns is to defend an externalist semantics against a raft of objections from advocates of a non-referential internalist semantics.

[19] For instance, the semantic minimalists Cappelen & Lepore (2005:148) see no problem in allowing speaker's intentions to play a role in fixing the semantic value of context-sensitive expressions, while Borg insists that semantic content is entirely formally driven (that is, it runs exhaustively off syntax and lexicon).

of indexical referents, he mentioned no other role for context or pragmatics in identifying that propositional content. It looks very much as if he too would have viewed the standing meaning (or 'conventional meaning', as he tended to phrase it) of a word as a fully semantic entity, hence a concept.

Unlike propositional minimalism, relevance theory is not bound to assume that words encode concepts. Given its basic premises that sentences do not encode thoughts/propositions and that pragmatic inference is essentially involved in the derivation of the explicature of an utterance, the way is open for word meanings to be something other than fully conceptual denotational entities.[20] In fact, as mentioned earlier, relevance theorists have argued that a sizable subset of the vocabulary encodes something other than concepts: 'procedures' in some cases (e.g. 'but', 'moreover', 'after all', 'anyway', 'well'); 'pro-concepts' in other cases (e.g. 'my', 'fast', 'near', 'do').[21] Nevertheless, Sperber and Wilson are quite clear that, in their view, the vast majority of open-class words encode concepts, which have a denotation, an externalist semantics, and so contribute directly to truth-conditional content. Here is a recent statement of that view: 'the ideas evoked by the presence of a word in an utterance are likely to be true of items in the linguistically specified denotation of the word, or, equivalently, of items in the extension of the concept encoded by the word.' (Sperber & Wilson 2008: 92). And this assumption about word meaning has underpinned much of my own work on lexical pragmatics within the relevance-theoretic framework (Carston 2002; Wilson & Carston 2007). In short, the standard RT position on these words is very similar to the minimalist semantic position: word meanings are concepts and concepts have a referential semantics (with the wrinkle mentioned in footnote 5).

So the overall picture of (encoded) word meanings in RT is quite heterogeneous: procedures of various sorts, pro-concepts, and concepts (denotational entities), all of which provide input to the pragmatic interpretation process, and all of which, concepts included, function merely as clues or evidence guiding and constraining processes of pragmatic inference whose goal is the recovery of the intended interpretation of the utterance. My suggestion is that the 'concepts' allegedly encoded by lexical forms are, in fact, conceptually underspecified meanings and that this meshes very well with the broader RT picture of lexical meaning and pragmatic processing. However, I want to keep these underspecified meanings distinct from pro-concepts, which have an indexical element in their lexical meaning. That is, they come with a parameter or parameters that must be given a contextual value, e.g. the meaning of 'my' probably includes a speaker parameter and a 'possession' parameter, the contextual filling of which, on a particular occasion of utter-

---

[20]   Again, the proviso here is that we need to give an account of how the move from the non-conceptual/non-semantic to the conceptual/semantic is made. I believe this can be done and without needing vast changes to the lexical pragmatic story told in section 2, but I am not attempting it in this paper.

[21]   There is work to be done in clarifying the differences between these two kinds of meaning. The key feature of pro-concepts is that they express concepts via an obligatory contextual specification of an indexical element contained in their lexical meaning. The inferential procedures associated with discourse connectives like 'however' are clearly distinct from these (see Blakemore 1987), but the term 'procedure' is now being used for a wider range of cases, including the meaning of pronouns, so it may be that pro-concepts are to be subsumed under this general label and that the taxonomic task ahead is to distinguish and characterise different subclasses of procedures.

ance of 'my book', could deliver a content of roughly the following sort: 'THE BOOK MARY WROTE', where Mary is the speaker of the utterance and the relevant 'possession' relation is the one between writer and thing written. The pragmatic process here is an instance of what Recanati (2004) calls variable saturation and it is mandatory: that is, in no matter what context, the speaker variable and the 'possession' relation variable must both receive a value if the utterance is to be fully comprehended (and if it is to be fully propositional). The underspecified meanings of open-class words (e.g. 'school', 'novel', 'milk', 'run', 'open', 'fly', 'red', 'circular', 'flat', 'raw') are different from these indexical words. First, the concepts they are used to express can differ from one another in arbitrarily many ways and their derivation depends on a wide array of different kinds of ordinary world knowledge which is unpredictable in advance of the utterance. Recall in this regard the different concepts of opening typically expressed by the following: 'open the door' (to let someone in), 'open the door' (to repair the lock), 'open the package', 'open your mouth', 'open your eyes', 'open someone else's mouth (or eyes)', 'open the iceberg', 'we'll have to open it' (which has many possibilities depending on the referent of 'it'), and so on, for everything that might, in some sense or other, be opened (and, of course, extra-linguistic context might well override the stereotypic or default mode of opening the object or entity denoted by the internal argument of 'open'). Second, even if it were possible to set out the full range of parameters of variance, it would not be mandatory (or even possible) to provide all of them with a semantic value on every occasion of use as only a subset would be contextually relevant. For more detail on the differences between indexicals and these cases of lexical meaning adjustment, see Recanati (2004: 94-95) and Bosch (2007, 2009).[22] I shall set aside indexicality and the notion of pro-concepts at this point and, for want of any better label, refer to the underspecified lexical meanings at issue here as 'concept schemas' or 'concept pointers'.

In a recent paper on the conceptual/procedural distinction, Deirdre Wilson (2011) has given much more weight and substance to the notion of procedural meaning in lexical semantics and has, in effect, suggested that all words encode a procedural component of meaning while some (the open classes) also encode a concept. Before looking more closely at this idea, let me say a little more about 'procedural' meaning as it has been developed in Relevance Theory. When it was first introduced (by Blakemore 1987), procedural meaning was taken to apply just to discourse markers, such as 'however', 'moreover', 'still', 'yet', 'but' and 'anyway', words that don't appear to contribute anything descriptive (truth-conditional) to the interpretation of an utterance. Blakemore proposed that what they do is provide a hearer with guidance on the kind of pragmatic inferences he is to perform in order to derive the intended implications (or implicatures) of an utterance. The notion of procedural meaning was extended in later work so as to encompass encoded constraints on other kinds of pragmatic tasks; for instance, it has been suggested that pro-

---

[22] Bosch (2009) distinguishes predicate indexicality from the kind of context-variable meaning of words that gives rise to polysemy. He discusses the adverb 'nearby' as a case of the former, the verb 'open' as a case of the latter, and the adjective 'fast' as a case where, as well as a lexically provided indexical element (for a comparison class), there is also this other kind of non-indexical context-dependence, which is not lexically specified but is wholly a matter of world knowledge and pragmatics. He provides some useful tests for distinguishing the two.

nouns encode procedural meaning which constrains the process of reference assignment, and that morphemes indicating grammatical moods, such as the indicative, the imperative, and the subjunctive, and modal particles (as in Japanese, for example) encode procedural meaning that constrains the pragmatic process of identifying the speaker's attitude or degree of commitment to the proposition she has expressed (Wilson & Sperber 1993; Wilson 2011). In short, the utility and appeal of the notion of procedural word meaning has grown considerably since it was first introduced, the central idea being that it is non-conceptual and constrains or guides pragmatic inference.

Let's move now to the new proposal in Wilson (2011), according to which 'open-class' or 'content' words encode both a concept and a procedure. She cites a communication from Dan Sperber in which he mentions the possibility that all words with a conceptual meaning also encode 'an instruction to inferentially construct an ad hoc concept using the encoded conceptual content as a starting point'. She endorses this suggestion and elaborates it as follows:

> 'On this approach, most words would encode some procedural content. Some would also encode conceptual content, whereas others (e.g. *however*) would not. Among words with both procedural and conceptual content, some (e.g. *giraffe*) would automatically trigger a procedure for constructing an ad hoc concept on the basis of the encoded concept, whereas others (e.g. *unless*) might encode a more specific procedure of the type familiar from Blakemore's work.' (Wilson 2011: 17).

She goes on to mention some advantages of this account over the standard RT position according to which these words encode just a concept. One is that it would make sense of the recurrent claim, arising from work in lexical pragmatics, that words function as 'pointers to' or 'pieces of evidence about' the speaker's meaning. Another is that it would dissolve a certain tension in the RT account of metaphorical and other nonliteral uses of words. The account has always rejected the Gricean treatment of nonliteral uses in terms of a flouting of a maxim of truthfulness and has maintained that it is not the case that the literal meaning (the encoded concept) is always the first to be considered as the correct interpretation and is only discarded in favour of another interpretation if it doesn't meet certain pragmatic standards (of informativeness, relevance, etc).[23] However, the worry is that, given that the relevance-based comprehension heuristic explicitly licenses hearers to follow a path of least effort in accessing and testing interpretations for relevance,[24] it seems natural to suppose that the encoded concept, which is made instantly available by the word form, would be tried first and only pragmatically adjusted if it didn't meet the required standards of relevance. The suggested move to incorporate into the meaning of content words a procedural component which requires that a relevance-driven process of concept construction is undertaken ensures that, although the encoded concept is activated

---

[23] There is quite a lot of experimental evidence now that disconfirms the 'literal first' hypothesis; for discussion, see Gibbs 1994, Giora 2003.

[24] In brief, the relevance-based comprehension heuristic says: (a) Follow a path of least effort in constructing an interpretation of the utterance; (b) Stop when your expectations of relevance are satisfied. For more detail, see Wilson & Sperber 2004, Sperber & Wilson 2012.

by the word uttered, it is not necessarily the first one to be composed into the interpretation.

This new conception of the meaning of open-class words as both conceptual and procedural raises a number of questions for me. First, it is difficult to see why a word that encodes a concept (a semantic entity with a 'linguistically specified denotation') would also encode a procedure that makes it obligatory for a hearer to build an ad hoc concept from the encoded one, especially when the encoded concept can, on occasion, be the concept communicated (Sperber & Wilson 1998, 2008). Second, the procedure involved would be identical across all the words which are taken to encode a concept, that is, the words 'giraffe', 'milk', 'money', 'run', 'pass', 'raw', 'red', and every other open-class word would all come with the same component of procedural meaning, namely, 'Construct an ad hoc concept based on the encoded concept', which seems odd since, by and large, the lexical meanings of words are distinct from each other and this goes as much for procedural meaning as for conceptual meaning, e.g. the procedural meaning of the pronouns 'he', 'she', 'it', 'we', 'they' is plainly distinct for each one, and linguists working on the procedural meaning of such closely related discourse connectives as 'but', 'however', 'nevertheless' and 'although' have put a lot of effort into pinpointing the fine differences in the inferential procedures they encode (see Blakemore 2002). Recognising the wide generality and recurrence of her proposed procedure, Wilson says that the procedure need not be attached to each word but rather could 'be formulated at the level of the class [of 'content' words] rather than the individual word' (2011: 18, n. 3). This might be feasible via some kind of lexical redundancy rule that says that, for every word that encodes a concept, there is a procedure to be followed of constructing an ad hoc concept on the basis of the given concept. Even so, it makes it very different from any of the more established elements of procedural meaning.

What really makes the proposal surprising to me, though, is that it seems entirely unnecessary on the relevance-theoretic account of utterance interpretation to issue instructions to the pragmatic system to construct ad hoc concepts. The goal of utterance interpretation is to recover a speaker's meaning, that is, the thought or thoughts she intends to communicate, where thoughts are structured arrays of concepts. On the account suggested, the words at issue encode concepts, so it is already evident that these words are contributors of concepts to the interpretation (as opposed to words like 'moreover' which do not encode anything conceptual). The general relevance-based comprehension heuristic takes care of the rest, that is, it ensures that the concepts recovered as speaker-meant are those that contribute to an optimally relevant interpretation, which may entail that the concept encoded is pragmatically adjusted (narrowed, broadened, or both), as discussed in section 2 and in greater detail in Wilson & Sperber (2002, 2004). So, the idea that, in addition to all this, every open-class word comes with (or triggers) an instruction to build an ad hoc concept seems like overkill.

However, I think there is a way of capturing the desirable aspects of the proposal while avoiding these problems and that is to construe word meaning along the lines I have been considering in this paper. The hypothesis is that so-called 'content' words have a semantically underspecified non-conceptual lexical meaning, that is, a concept schema (or conceptual pointer or indicator). On such an account, each word comes with its own dis-

tinct but schematic meaning, which functions as a constraint on the general pragmatic process of accessing or constructing a concept, a process which is wholly motivated by the goal of the pragmatic system which is to deliver speaker meaning. As the meaning provided by concept schemas is not fully conceptual, concept construction would be an obligatory pragmatic process and, like the various proposals for other kinds of word meaning, e.g. inferential procedures which constrain implicature derivation, indexicals which constrain reference assignment, pro-concepts which constrain pragmatic completion processes (e.g. 'my', 'fast'), these concept schemas or pointers would be linked to a particular pragmatic process, namely ad hoc concept construction. This sort of account, assuming it can be properly worked out, is not prey to the problems I mentioned above for the concept-plus-procedure account: it does not entail an obligatory process that is sometimes unnecessary (as when the encoded concept is the concept communicated), it doesn't entail a component of lexical meaning that is the same for thousands of words (that is, the instruction to build an ad hoc concept) and it doesn't formulate within the lexical semantics of a language a process (concept construction) that is entirely a matter of pragmatics. Furthermore, the advantages that Wilson discusses for the concept-plus-procedure account, are equally carried by this concept schema account: it makes perfect sense of the idea that all words are merely pointers to or evidence for a speaker's meaning, and, since there is no encoded concept (literal meaning), it allows for any one of a range of concepts to be the first one accessed or constructed, as determined by considerations of relevance (specifically, by the 'least effort' part of the relevance-based comprehension heuristic).

As discussed in sections 3 and 4, there are further advantages to the concept schema (or pointer) approach too. First, because words don't include full-fledged concepts on this sort of account, the issue of what exactly these concepts are, what they contribute to thoughts, what their denotations are, doesn't arise. Recall, in this regard, the question about the general concept encoded by the verb 'open', as distinct from all the more specific concepts it is used to express and communicate. Second, the schematic meaning view provides a basis for understanding polysemy (the fact that bundles of related senses seem to accrue around a word form), which is a very prevalent phenomenon of language and yet seems resistant to adequate explanation within many accounts of word meaning.[25] While Wilson's concept-plus-procedure hypothesis would certainly account for *pragmatic* polysemy (the construction in context of new senses and the retrieval of recurrent but not fully conventional senses), the concept schema account of lexical meaning does that equally well and it does more: by locating all polysemy (semantic and pragmatic) outside the lexicon (that is, in the conceptual system) and providing just a schematic, underspecified lexi-

---

[25] An approach like Pustejovsky's (1995) 'generative lexicon' has to import a great deal of general world knowledge into the lexicon and, even then, it succeeds in 'generating' only a small subset of polysemies (for insightful critiques, see Fodor & Lepore (1998) and Bosch (2009)). And the minimalist semantic approach doesn't seem to have the resources to distinguish polysemy from the different phenomenon of homonymy: word meanings on the minimalist approach are 'referential atoms', so it looks as if cases of highly conventionalised polysemy, for instance, words like 'school', 'newspaper' and 'disarm', discussed above in section 4, must be listed in the lexicon as a single lexical form with several distinct referential axioms, just as for homonyms like 'bank', 'coach' and 'bug'. In a sustained account of word meaning within semantic minimalism, Borg (2012) acknowledges the issues raised by polysemy (p. 172) and discusses and criticises existing approaches (pp. 177-89), but, as far as I can see, does not provide a minimalist solution.

cal meaning, it meshes neatly with experimental findings on the processing of polysemous words (Frisson 2009), as discussed in section 4.

Finally, I suggest that this non-conceptual account of word meaning accords better than an account that assumes words encode concepts (constituents of thoughts) with the following very recent reflection by Sperber and Wilson on the nature of linguistic meaning:

> '… as the gap between sentence meaning and speaker's meaning widens, it increasingly brings into question a basic assumption of much philosophy of language, that the semantics of sentences provides straightforward, direct access to the structure of human thoughts. We have argued that *linguistic meanings are mental representations that play a role at an intermediate stage* in the comprehension process. Unlike speakers' meanings (which they resemble in the way a skeleton resembles a body), linguistic meanings are not consciously entertained. In other words, whereas speakers' meanings are salient objects in personal psychology, *linguistic meanings only play a role in sub-personal cognition.*' (Sperber & Wilson 2012: 26-27; my emphases)

6. Conclusion: What isn't said and what cannot be said

In this paper, I've made a preliminary case for treating the standing or encoded meanings of open-class words as semantically underspecified (non-conceptual), thus as different in kind from the fully semantic entities that words are used by speakers to express or communicate, that is, concepts, components of the thoughts (propositions) expressed.

I will conclude with a brief reflection on the way in which this view of word meaning interacts with the wider linguistic underdeterminacy position and its implications for the semantic notion of 'what is said' (which consists of encoded or conventional sentence meaning with indexical references assigned). Among those who accept that sentence meaning typically underdetermines the proposition a speaker expresses and means (what she communicates explicitly), some take this to be an essential property of the relation between language and thought, while others take it to be merely a convenient effort-saving short-cut. According to the latter view, speakers could, by uttering more words, formulate their thoughts fully explicitly and thus 'say' what they mean, leaving nothing to the inferential pragmatic capacities of their addressees. For instance, Bach (1994) discusses a number of examples of sentences which when uttered require addressees to perform pragmatic processes of 'completion' (e.g. 'Mary is ready', 'He has taken enough') or 'expansion' (e.g. 'I haven't had breakfast', 'France is hexagonal'). According to him, '… what is being communicated could have been made *fully explicit* by the insertion of additional lexical material' (Bach 1994: 134). So a speaker could have uttered 'Mary is ready to leave for the airport' or 'France is approximately hexagonal' and thus succeeded in saying exactly what she meant. Speakers who choose to utter more abbreviatory sentences do so to save themselves some articulatory effort in situations where they can rely on their addressees to infer with ease the needed completion or expansion.

In more recent work, Bach goes further in denying that there is any significant difference between thoughts and the meanings of natural language sentences: '… it is arguable that all of our thoughts are explicitly expressible, in which case for every thought there is at least one sentence that would express it explicitly.' (Bach 2010: 129). What this entails is that, as well as a lot of semantically incomplete gappy sentences, useful for speedy conversational exchanges, public languages provide sentences which fully express thoughts, and, in fact, are rich enough to do this for all and any thoughts we may have. This is a very strong version of a principle of expressibility or effability, according to which sentences fully encode thoughts, modulo a small set of indexical words ('I', 'you', 'today', 'now', 'here') whose content is fixed by a few parameters of semantic context (speaker, hearer, day of utterance, place of utterance), independent of any pragmatic considerations. So while speakers may often mean more than they say and pragmatics may be needed to complete and expand what they said in order to arrive at the proposition they intended, they could, assuming they are fully linguistically competent, have found a sentence the utterance of which would have expressed their thought fully and so required no pragmatic work on the part of their addressee. In short, language is such that, if they so chose, competent speakers could always say exactly what they mean.

Relevance theorists (e.g. Carston 2002, Sperber and Wilson 1986/95, 2012) and contextualists (e.g. Recanati 1994, 2004) have taken a very different view of linguistic underdeterminacy, which is that it is essential, that is, that there are always components of a speaker's meaning which the linguistic expressions she employs do not encode. Although the addition of more words may make a speaker's primary meaning (her 'explicature') more explicit, full explicitness (full encoding of content) is quite generally not achievable. I won't rehearse the arguments for this here (see the references just given), but simply point out that an account of word meanings as nonconceptual (semantically underspecified) would be the completing component of this view of the relation between language and thought: not only do sentence meanings underdetermine thoughts, but the basic constituents of sentences (words) underdetermine the basic constituents of thoughts (concepts). If this account turns out to be right, it's not just that we don't always say what we intend our hearers to take us to mean but that it is simply not possible to say what we mean.

References

Bach, K. 1994. Conversational Impliciture. *Mind and Language* 9: 124-65.

Bach, K. 2005. Context *Ex Machina*. *Semantics Versus Pragmatics*, ed. Z. Szabo, 15-44. Oxford: Clarendon Press.

Bach, K. 2010. Impliciture vs Explicature: What's the difference? *Explicit Communication: Robyn Carston's Pragmatics*, eds. B. Soria and E. Romero, 126-137. Basingstoke: Palgrave.

Blakemore, D. 1987. *Semantic Constraints on Relevance*. Oxford: Blackwell.

Blakemore, D. 2002. *Relevance and Linguistic Meaning: The Semantics and Pragmatics of Discourse Markers*. Cambridge: Cambridge University Press.

Bloom, P. 2000. *How Children Learn the Meaning of Words*. Cambridge, Mass.: MIT Press.

Borg, E. 2012. *Pursuing Meaning*. Oxford: Oxford University Press.

Bosch, P. 2007. Productivity, Polysemy, and Predicate Indexicality. *Proceedings of the Sixth International Tbilisi Symposium on Language, Logic & Computation*, eds. B. ten Cate and H. Zeevat, 58-71. Berlin: Springer.

Bosch, P. 2009. Predicate Indexicality and Context Dependence. *Utterance Interpretation and Cognitive Models: Current Research in the Semantics/Pragmatics Interface vol. 20*, eds. P. De Brabanter & M. Kissine, 99-126. Emerald Group.

Cappelen, H. and E. Lepore. 2005. *Insensitive Semantics. A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell

Carston, R. 2002. *Thoughts and Utterances: the Pragmatics of Explicit Communication*. Oxford: Blackwell.

Carston, R. 2004. Relevance Theory and the Saying/Implicating Distinction. *The Handbook of Pragmatics*, eds. L. Horn and G. Ward, 633-56. Oxford: Blackwell.

Carston, R. 2009. Relevance Theory: Contextualism or Pragmaticism? *UCL Working Papers in Linguistics* 21: 17-24.

Carston, R. 2010. Metaphor: *Ad Hoc* Concepts, Literal Meaning and Mental Images. *Proceedings of the Aristotelian Society* 10: 295-321.

Carston, R. 2012. Word Meaning and Concept Expressed. *The Linguistic Review* 29: 607-23.

Carston, R. and A. Hall. 2012. Implicature and Explicature. *Cognitive Pragmatics,* vol.4 of *Handbooks in Pragmatics*, eds. H-J. Schmid and D. Geeraerts, 7-84. Berlin: Mouton de Gruyter.

Carston, R. and C. Wearing. 2011. Metaphor, Hyperbole and Simile: A Pragmatic Approach. *Language and Cognition* 3: 283-312.

Fodor, J. A. 1998. *Concepts. Where Cognitive Science Went Wrong.* Oxford: Clarendon Press.

Fodor, J. A. 2001. Language, Thought and Compositionality. *Mind and Language* 16: 1-15.

Fodor, J. A. and E. Lepore. 1998. The Emptiness of the Lexicon: Reflections on Pustejovsky. *Linguistic Inquiry* 29 (2): 269-88. Reprinted in Fodor and Lepore 2002: 89-119.

Fodor, J. A. and E. Lepore. 2002. *The Compositionality Papers*. Oxford: Clarendon Press.

Frazier, L. and K. Rayner. 1990. Taking on Semantic Commitments: Processing Multiple Meanings vs. Multiple Senses. *Journal of Memory and Language* 29: 181-200.

Frisson, S. 2009. Semantic Underspecification in Language Processing. *Language and Linguistics Compass* 3: 111-127.

Frisson, S. and M. Pickering 2001. Obtaining a Figurative Interpretation of a Word: Support for Underspecification. *Metaphor and Symbol* 16: 149-171.

Gibbs, R. 1994. *The Poetics of Mind: Figurative Thoughts, Language and Understanding*. Cambridge: Cambridge University Press.

Giora, R. 2003. *On Our Mind: Salience, Context and Figurative Language*. New York: Oxford University Press.

Grice, H. P. 1967/89. *Logic and Conversation*. William James lectures, 1967. Harvard University. Printed in Grice, H. P. 1989: 1-143.

Grice, H. P. 1975. Logic and Conversation. *Syntax and Semantics 3: Speech Acts*, eds. P. Cole and J. Morgan, 41-58. New York: Academic Press. Reprinted in Grice, H.P. 1989: 22-40.

Grice, H. P. 1978. Further Notes on Logic and Conversation. *Syntax and Semantics 9: Pragmatics*, ed. P. Cole, 113-127. New York: Academic Press. Reprinted in Grice, H.P. 1989: 41-57.

Grice, H. P. 1989. *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.

Hampton, J. 2011. Concepts and Natural Language. *Concepts and Fuzzy Logic*, eds. R.Belohlavek and G.J.Klir, 233-258. Cambridge Mass.: MIT Press

Horsey, R. 2006. *The Content and Acquisition of Lexical Concepts*. PhD Dissertation, University College London.

Kempson, R., E. Gregoromichelaki and C. Howes. 2011. Introduction. *The Dynamics of Lexical Interface*s, eds. R. Kempson, E. Gregoromichelaki & C. Howes, 1-20. Stanford: CSLI Publications.

Lieber, R. 2004. *Morphology and Lexical Semantics.* Cambridge: Cambridge University Press.

Murphy, G. L. 2002. *The Big Book of Concepts*. Cambridge, Mass.: MIT Press.

Pickering, M. and S. Frisson 2001. Processing Ambiguous Verbs: Evidence from Eye Movements. *Journal of Experimental Psychology: Learning, Memory and Cognition* 27: 556-73.

Pritchard, T. forthcoming. Affordance, Detachment, and our Understanding of 'Pencil'.

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

Recanati, F. 1994. Contextualism and Anti-Contextualism in the Philosophy of Language. *Foundations of Speech Act Theory*, ed. S. Tsohatzidis, 156-66. London and New York: Routledge.

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

Ruhl, C. 1989. *On Monosemy: A Study in Linguistic Semantics*. Albany: State University of New York Press.

Sperber, D. and D. Wilson. 1986/95. *Relevance: Communication and Cognition*. Oxford: Blackwell; Cambridge, Mass.: Harvard University Press. Second edition (with postface) 1995.

Sperber, D. and D. Wilson. 1998. The Mapping Between the Mental and the Public Lexicon. *Language and Thought*, eds. P. Carruthers and J. Boucher, 184-200. Cambridge: Cambridge University Press. Reprinted in Wilson and Sperber 2012: 31-46.

Sperber, D. and D. Wilson. 2008. A Deflationary Account of Metaphors. *The Cambridge Handbook of Metaphor and Thought*, ed. R. Gibbs, 84-105. Reprinted in Wilson and Sperber 2012: 97-122.

Sperber, D. and D. Wilson 2012. Introduction: Pragmatics. In Wilson and Sperber 2012: 1-27.

Travis, C. 2000. *Unshadowed Thought*. Cambridge Mass.: Harvard University Press.

Tyler A. and V. Evans 2001. Reconsidering Prepositional Polysemy Networks: The Case of 'over'. *Language* 77: 7214-65.

Vega Moreno, R. 2007. *Creativity and Convention: the Pragmatics of Everyday Figurative Speech*. Amsterdam: John Benjamins.

Wharton, T. 2002. Paul Grice, Saying and Meaning. *UCL Working Papers in Linguistics* 14: 207-48.

Wilson, D. 2003. Relevance Theory and Lexical Pragmatics. *Italian Journal of Linguistics/Rivista di Linguistica* 15: 273-291.

Wilson, D. 2011. The Conceptual-Procedural Distinction: Past, Present and Future. *Procedural Meaning: Problems and Perspectives*, eds. V. Escandell-Vidal, M. Leonetti and A. Ahern, 3-31. Emerald Group Publishing.

Wilson, D. and R. Carston. 2006. Metaphor, Relevance and the 'Emergent Property' Issue. *Mind and Language* 21: 404-33.

Wilson, D. and R. Carston. 2007. A Unitary Approach to Lexical Pragmatics. *Pragmatics*, ed. N. Burton-Roberts, 230-60. Basingstoke: Palgrave.

Wilson, D. and D. Sperber. 1993. Linguistic Form and Relevance. *Lingua* 90: 1–25.

Wilson, D. and D. Sperber. 2002. Truthfulness and Relevance. *Mind* 111: 583-632.

Wilson, D. and D. Sperber. 2004. Relevance Theory. *The Handbook of Pragmatics*, eds. L. Horn and G. Ward, 607-32. Oxford: Blackwell.

Wilson, D. and D. Sperber. 2012. *Meaning and Relevance*. Cambridge: Cambridge University Press.