

Word meaning and concept expressed¹

ROBYN CARSTON

Abstract

The concept expressed by the use of a word in a context often diverges from its lexically encoded context-independent meaning: it may be more specific or more general (or a combination of both) than the lexical meaning. Grasping the intended concept involves a pragmatic process of relevance-driven adjustment or modulation of the lexical meaning in interaction with the rest of the utterance and with contextual information. The issue addressed here is the nature of the input to the pragmatic process of meaning adjustment, that is, the nature of the standing (encoded) meaning of the word type. The widespread assumption that lexical meaning is conceptual, hence directly expressible, is challenged and a case made for the merits of an account of word type meaning in non-conceptual terms.

1. What is word meaning?

We use sentences to express/communicate thoughts (truth-conditional contents) and we use words to express/communicate concepts, which are constituents of thoughts (hence contribute to truth-conditional contents). It is now quite widely accepted that the meaning (or semantic content) that a word is used to express or communicate on an occasion of utterance is often distinct from the meaning it has as an expression type in a language system (that is, its standing or encoded meaning). This view is shared by 'contextualist' philosophers of language, by some linguists, and by pragmaticists working within the cognitive-scientific framework of relevance theory. The main aim of this paper is to consider the nature of the context-free word meaning which is the starting point for the pragmatic processes that deliver the occasion-specific meaning (the concept meant or communicatively intended by the speaker).

I am confining my attention here to what are often called 'open class' words, that is, words with an apparently descriptive meaning, such as nouns, verbs, adjectives, and most adverbs, leaving aside 'closed class' words, such as indexicals, determiners, function words, and connectives. With the domain so restricted, the question is: 'what is a word meaning?' Word type meanings might be concepts, hence contentful entities that can be constituents of thoughts. If that is the right answer, then standing word meanings are the same sort of thing as that which we use them to express/communicate (concepts, semantically

¹ I thank Mark Textor, Ruth Kempson and Timothy Pritchard for very helpful discussions. I am also grateful to two anonymous referees for *Linguistic Review* and, in particular, to the editor of this special issue, Stavros Assimakopoulos, for his sound advice and constant encouragement. This work is supported by an AHRC Research Grant, AH/I000216/1, awarded to the project: *Word Meaning: What it is and what it is not.*

evaluable entities) and there is no reason why speakers would not, at least sometimes, use words to express the very meaning they encode. This view is held by some psychologists (e.g. Murphy 2002), by some philosophers (e.g. Fodor 1998) and by relevance theorists (e.g. Sperber and Wilson 1986/1995, 1998). However, there is an equally widely held view that word meanings are 'underspecified'; that is, that they cannot contribute directly, without modification or transformation of some sort, to the thoughts/propositions that utterances in which they occur are used to express. There are various possibilities for what these underspecified entities might be: a special kind of 'lexical' concept, a pro-concept, a schema or procedure or set of constraints on the kind of contentful concept they can be used to express/communicate. Any of these could qualify as a word meaning, as a specification of a word's 'semantic potential'. However, a more radical position is that words (lexical forms) do not encode concepts or abstract schemas or constraints, but are associated with something else altogether, something that does not qualify as a meaning of the expression type. Two apparently rather different possibilities that have been suggested are (a) collections of memory traces or exemplars of previous uses (tokenings) and (b) bundles of contingent encyclopaedic information about the things in the world the word is used to refer to.

These four positions (ordinary concepts, lexical concepts, abstract schemas or constraints, previous uses or encyclopaedic information) can all be seen as falling within a broadly contextualist view of word meaning and they correlate roughly with the four contextualist positions set out by Recanati (2004). The most conservative of these is the 'strong optionality' position (quasi-contextualism), according to which a word's meaning may contribute directly (unmodulated) to truth-conditional content *or* may be pragmatically modulated/adjusted. The second is the 'pragmatic composition' view, according to which a word meaning (lexical concept) could be an expressed sense, but the process of composing it together with the other words in the utterance forces its pragmatic adjustment. On the third, more radical, view which Recanati calls the 'wrong format' position, word meanings cannot enter directly into thought, but must be transformed into the right (conceptual/contentful) format, presumably by some pragmatic interpretive process or other. Here, there is a distinction between the 'semantic structures' of linguistic expressions and the conceptual structures of thought. If word meanings are abstract schemas or sets of constraint or rules for use, then they are in the 'wrong format' to be semantically contentful or to be constituents of thoughts. Finally, there is the most extreme position, that of meaning eliminativism, according to which words (qua types) do not have meanings at all; only tokens (specific utterings) of words have meanings.

It is instantiations of these latter two broad positions that I want to consider as possible candidates for what word types bring to the pragmatic processes of utterance understanding. I will present a range of considerations that point in the direction of there being something less than a fully conceptual meaning for words, thus ruling out the first two positions. I have no knock-down arguments and there are some obvious pitfalls to taking on this non-conceptual view. However, I do think it merits serious consideration, especially when placed in the context of relevance-theoretic lexical pragmatics, which offers an

explanatory account of how the concept a speaker expresses with a word can vary across occasions of use.²

2. Relevance theory pragmatics

The view that sentence meaning (logical form) is seldom, if ever, fully propositional has been a basic tenet of relevance theory since its inception. The claim is that it is a propositional template or radical, which must be pragmatically completed and elaborated on each occasion of utterance in order to derive the propositional content meant (communicatively intended) by the speaker (Carston 2002). Recently, the focus has been more on word meaning and a detailed relevance-theoretic account of lexical pragmatics has been developed (Wilson and Carston 2007). Consider the examples in (1), a set of utterances about Boris, who, let us suppose, is a man in his forties and has been married for many years:³

- (1) a. *Boris is a man.*
 Encoded meaning: B IS A MAN
 Explicature: B IS A MAN*
b. *Boris is a child.* CHILD*
c. *Boris is a bachelor.* BACHELOR*
d. *Boris is a chameleon.* CHAMELEON*

In (1a), the linguistically encoded content is a trivial truth, hence uninformative, insufficiently relevant. The addressee's process of trying to derive contextual implications from the utterance, in accordance with his expectation of relevance, leads to the encoded concept MAN being narrowed down so as to encompass just men of some kind. Depending on the specifics of the context, it could be narrowed down to 'typical man' or 'ideal man' and, of course, what constitutes a typical man or an ideal man will itself vary from context to context. The outcome of this process is an occasion-specific sense (or 'ad hoc' concept) MAN*, which picks out a proper subset of the set of individuals that fall under the original encoded concept MAN.

In (1b), we have the opposite phenomenon: the encoded concept CHILD is adjusted so as to mean roughly 'person who behaves in certain childish (or child-like) ways', and the result is a concept CHILD* which is broader than the lexically encoded concept - it includes actual children and some adults. Then, if we take (1c) as an utterance by Boris's wife, who has long endured his affairs with other

² An anonymous referee has suggested that I should make it explicit early on that, in this paper, I am treating all (open class) words as having the same kind of meaning and not discriminating between those words that are (semantically) polysemous and those that are not. In fact, I am far from sure that any such distinction should be made: my working hypothesis is that (open class) words quite generally are susceptible to pragmatic adjustment in context and that some, a minority, of these (initially ad hoc) derived senses become routinized or conventionalized to varying degrees, due to repeated use. What we describe as 'polysemous words' are those whose several senses have crossed some threshold of frequency or conventionality.

³ The '*' on the concept expressed by the speaker of these utterances is just a notational device to indicate that it is distinct from the lexically encoded concept and was derived by a pragmatic process.

women and general lack of commitment, this is, arguably, both a broadening of the lexical concept *BACHELOR* (it includes *married* men who behave in certain ways) and a narrowing (it excludes bachelors who don't behave in this stereotypical way). Finally, (1d) is a typical metaphorical use, which in standard relevance theory is taken to be a radical kind of broadening, so the concept communicated, *CHAMELEON**, roughly paraphrased as 'individual that can alter appearance or behaviour so as to fit in with current surroundings or circumstances', includes actual chameleons (a kind of lizard) and certain human beings (and any other creatures with the property at issue).⁴

These pragmatically-derived (ad hoc) concepts are components of the speaker's explicature, or what Recanati (2004) speaks of as the pragmatically enriched 'what is said', the intuitive truth-conditional content of the utterance. They are recovered by the addressee in the process of finding an interpretation that meets his *context-specific expectations of relevance* (licensed by the general presumption of 'optimal relevance' carried by all utterances).⁵ According to the account, there is a *single process* of lexical concept adjustment (or meaning modulation), which can have any of *several outcomes*: it can result in a concept whose denotation is narrower or broader (or both) than that of the lexical concept.⁶ The process is a kind of 'free' pragmatic enrichment; that is, unlike the pragmatic fixing of a value for an indexical, it is not linguistically mandated or controlled.

Now, let's consider what it is that words, qua expression types in an individual's language system, bring to the process of interpreting particular utterances; that is, what the lexically-based input to pragmatics is, according to the relevance-theoretic approach (Sperber and Wilson 1998; Wilson and Carston 2007). In effect, there are two parts to this, a semantic part, the encoded meaning of the word, and a contingent, non-semantic part, consisting of encyclopaedic information associated with the encoded meaning. Let's look at these in turn.

Most words are taken to encode concepts, specifically atomic concepts. There is a simple mapping from lexical form, e.g. /kæt/, to mental concept, e.g. CAT, the concept is unstructured and the lexical entry does not specify any further information about its content or semantic behaviour. In short, the position is essentially the same as Jerry Fodor's 'disquotational lexicon': the word 'house' means HOUSE, 'miserable' means MISERABLE, 'keep' means KEEP, etc.

⁴ An anonymous referee has objected to this account of the metaphorical use of the word 'chameleon', maintaining that the concept that it expresses pertains just to a certain kind of human being (so would not include actual chameleons in its denotation). I am broadly sympathetic to this position and am currently developing an account of metaphorical language use according to which it requires both broadening and narrowing, so that the denotation of the ad hoc concept derived may merely overlap with that of the lexical concept or may be entirely disjoint from it (see Carston and Wearing [2011] for a preliminary formulation of this idea). However, the position outlined above is the established 'loose use' account of orthodox relevance theory (see, for instance, Sperber and Wilson [1986/1995]).

⁵ For a recent outline of the principles and mechanisms of relevance theory, see Carston 2012.

⁶ Several people have raised objections to the construal of concept broadening and narrowing in denotational terms (Kempson [p.c.] and Textor and Allott [forthcoming]). I see this externalist semantic way of characterising these notions as a first shot, made in the interests of presentational clarity, but as only a part (and a part that may need to be modified) of a much more extensive account of ad hoc concepts whose internalist mental representational details must ultimately be given.

(Fodor 1998; Fodor and Lepore 1998), and most of the details of Fodor's view of concepts are preserved: Concepts are mental particulars (they function as mental causes and effects); they categorise the world so applications of concepts are susceptible of 'semantic evaluation' (as true/false; correct/incorrect); concept content is wholly referential (not constituted by inferential relations);⁷ concepts are constituents of thoughts and of each other (phrasal concepts); thoughts (and phrasal concepts) inherit their content from the contents of their constituent concepts. In a sense, then, words (lexical forms) don't have a semantics – they merely inherit a semantics from the concepts they encode – and the 'real' semantic story is about the concept-property (mind-world) relation, hence what it is to possess a concept and how we come to possess them, that is, what the mechanisms are through which a symbol in the head locks onto a property in the external world.

Associated with concepts (whether lexicalised or not) are collections of encyclopaedic information, including general knowledge and individual beliefs about the things they denote, 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.) As Fodor (2008) puts it, concepts can be thought of as names or labels for files containing such collections of contingent information, so that: 'In effect, according to this story, *we think in file names*;' (Fodor 2008: 95).

Encyclopaedic information plays a key role in the relevance-theoretic process of lexical concept adjustment. 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, concepts in the decoded logical form are adjusted by backwards inference, so that only implications that are ultimately grounded in the explicature are confirmed. The overall interpretation is accepted provided it meets the addressee's expectation of relevance. Consider again example (1b) '*Boris is a child*'. Depending on the wider discourse situation, contextual implications such as Boris is sweet and innocent, untouched by life experience, may be inferred, based on assumptions accessed from the encyclopaedic entry for CHILD, which, by backwards inference, lead to a particular ad hoc concept CHILD*. In another utterance situation, different items of encyclopaedic information about children might be more highly activated making most accessible such implications as that Boris doesn't earn his keep, expects others to look after him, is irresponsible, etc., resulting in a distinct ad

⁷ Here relevance theory does depart from Fodor in assuming that many concepts come with a logical entry, a set of inference rules that capture certain necessary conditions on a concept's content. I can and will ignore this complication for the purposes of this paper.

hoc concept CHILD** in the explicature. And there are other – perhaps indefinitely many – possibilities.

There's no arguing with the existence and importance of this kind of general world knowledge associated with words and concepts, nor with the claim that it is contingent, that is, it is extrinsic to the concept's content and plays no part in its individuating conditions. The part of the story that one could take issue with is the linguistic semantic part; that is, the claim that word type meanings are concepts. In the next section, I'll set out some considerations that make an alternative non-conceptual account worth serious investigation.

3. Considerations in favour of a schematic (underspecified) lexical meaning

I take it that concepts are, first and foremost, constituents of thoughts. This is in accordance with Fodor's view of them as, in effect, words of Mentalese (the language of thought) and with the views of most people working within cognitive science, including relevance theorists. So the first point I want to consider here concerns the role of the hypothesised encoded word meanings (lexicalised concepts) as components of thought. Take the concept HAPPY, for instance, (allegedly) encoded by the word 'happy'; this provides communicative access to a wide range of other more specific concepts, including one for a steady state of low-key well-being (HAPPY*), another for a momentary experience of intense joy (HAPPY**), another for the sense of satisfaction that accompanies a successful transaction or completion of a job (HAPPY***), and so on. All of these specific concepts are components of thoughts we might have. The question then is: Does the very general lexically-encoded concept HAPPY occur as a component of thought and, if so, what sort of thought is this?

Similar questions arise for the (alleged) concept OPEN encoded by the verb 'open'. In an interesting discussion of pragmatic polysemy, Sperber and Wilson (1998: 197) say: "A verb like 'open' acts as a pointer to indefinitely many notions or concepts ..." and they mention cases for which the intended concept is jointly indicated by the verb and its direct object, as in '*open the door*', '*open a letter*', '*open a tin*', '*open one's mouth/one's eyes*', etc, and others which depend on broader non-linguistic context (e.g. when opening a door might involve breaking it down with an axe, or opening one's mouth might require removal of stitches with which it has been fastened). However, they maintain that the pragmatic adjustment process is optional because: "It may so happen that the intended concept is the very one encoded by the word, which is therefore used in its strictly literal sense" (Sperber and Wilson 1998: 196-197). So let's consider the assumed lexical concept OPEN and what it is to have 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*', etc. The question is whether there is any definite thought at all or whether any thought about opening must contain one of the more specific concepts. Consider the following, trying to construe them as thoughts containing a very general lexical concept OPEN as a constituent:

- (2) a. *Whenever I open anything I feel anxious.*

b. *Everyone opens things sometimes.*

As far as I can tell, the thoughts about ‘opening things’ that we take to be expressed by these sentences are ones in which the ‘things’ at issue are construed as some sort of coherent subcategory of all the things that one could talk about opening, so, e.g., it might be the category of things that can contain stuff inside them, like boxes, envelopes, files, brief-cases, and cupboards. It seems unlikely to include opening one’s mouth or eyes (or even opening curtains or windows, or gates), let alone opening discussions, lectures, issues, minds, hearts or cans of worms. So while the ‘open’ concept that will figure in the thoughts in (2) is indeed quite a broad one, it is still, I think, considerably narrower than the supposed lexical concept that is allegedly encoded by the verb ‘open’ and provides the basis for inferring all the more particular concepts of opening.

This first consideration has been presented in the form of a thought experiment and it would, of course, be more satisfactory if it could be backed up by more empirically-based evidence. This might be possible through behavioural experiments on human categorization, which, if the thinking here is right, employs concepts of a finer-grain than such very general lexically encoded concepts as OPEN and HAPPY.

The second consideration hinges on some well-aired problems with the phenomenon of polysemy. Let’s take a fresh example, the word ‘stop’, as discussed by Agustín Rayo, who notes that it can move across syntactic categories (verbal and nominal):

You can stop writing; you can stop a burglar; you can stop a cheque; you can stop a nail hole with plaster; you can use your fingers to stop the holes of a flute; you can stop a poker into the fire; you can stop the tide by anchoring your boat. ... You can come to a stop; there can be a stop in your speech; you can include a stop in a telegram; you can put a stop on a camera; you can pull out all the stops on an organ. (Rayo forthcoming)

The language is full of such polysemous verbal-nominal words; consider, for instance, ‘fall’, ‘rest’, ‘cut’, ‘run’, ‘jump’, ‘skip’, ‘walk’, ‘start’, ‘end’, ‘turn’, ‘slip’, ‘pass’, ‘talk’, ‘sign’, ‘file’, and so on.⁸ As Rayo points out, this kind of grammatical and semantic versatility has been tackled by computationally-minded linguists through the postulation of an array of (sometimes quite complex) lexical-semantic rules (see, e.g., Pustejovsky 1995, Asher 2011).

There is a major obstacle to any lexical rule approach which is that since the range of concepts that a word can be used to convey is indefinite, commensurate with the indefinite range of contexts human communicators can find themselves in, a line has to be drawn between those that are to be accounted for by linguistic rules or conventions and those that are to be left to the pragmatic ingenuity of language users. But, as Nunberg pointed out long ago: “There is a substantial class of cases where we have no principled grounds for

⁸ On the account of the (‘exo-skeletal’) lexicon developed by the syntactician Hagit Borer, words as isolated entities (or ‘listemes’, as she calls them) do not belong to specific syntactic categories but only acquire syntactic status when they are used in particular syntactic structures (see Borer 2005). This view appears likely to mesh better with the ideas being pursued here than the predominant view among syntacticians which has much of the grammar projected from complex lexical entries (the ‘endo-skeletal’ view of the lexicon).

deciding which of several uses is conventional, i.e. licensed entirely by linguistics rules, [and which are derived]" (Nunberg 1979:154). The most problematic manifestation of this issue is what he called 'the non-uniqueness of semantic solutions', that is, cases for which there does not seem to be, among the related senses, one that is the basic or central sense, from which the others are derived.

Consider the following senses of the word 'window':

- (3)
- a. *The bay windows are a beautiful feature of the house.* [glass pane and frame]
 - b. *The cricket ball smashed my study window.* [glass pane]
 - c. *She crawled through the upstairs window and fell onto the floor.*
[open space in wall]
 - d. *The eyes are the window to the soul.* [something one can see through]
 - e. *We must seize this window of opportunity.* [something of short duration]

Which (of the first three) of these is the lexical concept WINDOW from which the others are pragmatically derived? None seems more intuitively basic than the others, nor the most useful as a starting point from which the others could be pragmatically derived in context. Cases of part-whole or metonymically related senses, which are surely among the most ordinary, everyday and uncreative of our uses of language, seem to be highly problematic in this regard. The noun 'novel', for example, discussed at length by 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). Again, 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: 59). And this point applies to a wide range of other words (see Nunberg 1979, Bosch 2007, Lossius Falkum 2011). So, instead of trying to force one of the multiple senses into the role of basic underived sense, perhaps we should give up on assuming that there is one that plays this role, that is, give up on assuming that there is an encoded lexical concept.⁹

A third consideration, one that has been pointed out by Sperber and Wilson (1998), is that 'words behave as if they don't encode concepts'. First, they suggest that there are many words that do not encode a full-fledged concept but what might be called a pro-concept, giving as likely examples, 'my', 'have', 'near', 'long', and saying that "while each of these examples may be contentious, the existence of the general category should not be" (Sperber and Wilson 1998: 185). They don't say anything more about what a pro-concept is, but it's clearly intended to be something less than a complete concept with a referential content, something that requires that a semantic value is pragmatically inferred in context, so it seems to be an indexical element of some kind. They then go on to say: "... quite commonly, all words behave *as if* they encoded pro-concepts:

⁹ Space limitations preclude discussion of the recent sophisticated and comprehensive account of word meaning by Asher (2011), in which certain polysemous words are taken to encode a complex semantic type, e.g. the lexical semantics of the word 'book' is of type [PHYSICAL OBJECT . INFORMATION], which is one of several complex types [α . β].

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" (Sperber and Wilson 1998: 185). And, in later, more fully developed work on lexical pragmatics, Wilson and Carston (2007: 231) maintain: "... 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*." So the question here is: if words quite generally behave as if they don't encode concepts, why maintain that they do encode concepts?

This is not just an issue for relevance theorists; it arises also in interestingly parallel recent work by the formally-oriented linguist Peter Bosch (2007). He distinguishes what he calls 'lexical concepts' from 'contextual concepts' and says: "cases of apparent variation in word sense require treatment at the conceptual level rather than a lexical semantic solution" (Bosch 2007: 58). He discusses a range of cases of nouns, including 'novel', as discussed above, and of predicates, including 'is working':

- (4) a. *Where is Fred?* *He's working.*
b. *What is Fred up to today?* *He's working.*
c. *How can Fred afford those expensive holidays?* *He's working.*
d. *Fred is working.*

As an answer to the different questions in (4a) - (4c), the utterance of '*He's working*' is interpreted differently in each case: as giving information about Fred's location (4a), about Fred's current activity (4b), and about Fred's financial situation (4c). Then, in (4d), it is interpreted differently depending on whether Fred is taken to be the name of our building's caretaker or of a prize race-horse or of my computer. As Bosch points out, these different contextual concepts expressed by the word 'work', WORK₁, WORK₂, WORK₃, WORK₄ ..., are truth-conditionally relevant and licence different inferences. He maintains that the lexical type meaning of 'work' and of many other words is 'underspecified', that is, it must be developed at a conceptual or pragmatic level in order for its expressed meaning to be realised. However, he persists in labelling it a 'lexical concept' and likens his position to that of Fodor and Lepore's (1998) 'disquotational' view of the lexicon (lexical forms map directly to atomic concepts). Again, my question is: why insist that words encode concepts?

The fourth and final consideration concerns semantic compositionality. As Fodor has pointed out repeatedly, there are some fundamental properties of language and thought, namely their systematicity and productivity, that can only be explained by the semantic compositionality of these representational systems: there is a basic stock of primitives (words/concepts) with stable semantic values and a recursive syntax such that the semantic value (content) of any sentence/thought is a function of the semantic value of the primitives and the way in which they are syntactically combined. With regard to the 'compositionality' of thoughts/concepts, Fodor (1998: 25-27) says: "Since it's required to explain productivity and systematicity, compositionality is, as one says, 'not negotiable'. An account of concept possession that is incompatible with the compositionality of thought is, ipso facto, out of the running." Thus, he has argued that an account of concept content in terms of stereotypes or prototypes

or inferential roles or partial definitions fails because these entities do not meet the compositionality requirement.

Now it might seem that this carries over point for point to public language systems and so to word meanings. However, that is not the case. Natural language sentences are simply not compositional in the required sense, that is, the propositions/thoughts they express are not determined by word type meanings and syntax alone, 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 doesn't express the thought that it would if the sentence were compositional. Either (the typical case) it *vastly underdetermines* the right thought; or the thought it determines when compositionally construed isn't, in fact, the one that it conventionally expresses." (Fodor 2001: 12, emphasis mine). He concludes: "The evidence suggests strongly that language is not compositional." (Fodor 2001: 14).

As mentioned in Section 2, this has always been a central claim of relevance theory and of the contextualist philosophers. What has not been so much noted is the implication that it has for an account of word (type) meanings: the compositionality-based arguments against the adequacy of prototypes, inferential roles, partial definitions, etc. as the semantic content of concepts do not carry over to word meanings. Once you drop the compositionality requirement on linguistic (sentence) meaning (while of course maintaining it for thoughts), any requirement that a word (type) meaning contributes content directly to truth conditions and so encodes a concept falls away. Words could encode prototypes, or inferential relations, or schemas, or constraints, or any of a range of other non-conceptual (non-contentful) possibilities, provided that the kind of component required to preserve the compositionality of thought, that is, concepts (with a referential semantics), can be delivered by pragmatic processes. In short, severing the relation between word type meanings and concepts does not violate the non-negotiable compositionality constraint.

I hope that collectively these considerations provide sufficient impetus to warrant looking into non-conceptual characterisations of word type meaning.

4. Alternatives to concepts as word 'meanings'

The four alternatives I'm going to mention briefly here are either instantiations of the position Recanati (2004) calls the 'wrong format' view, i.e. linguistic 'semantics' does not provide a truth-conditional component (a 'content'), but still there is some stable, context-free *meaning* associated with the word type, or instantiations of the most radical contextualist position, which he calls 'meaning eliminativism', i.e. there is nothing resembling a stable word type meaning, but rather a collection of resources for concept-making that the lexical form activates.

The first possibility to consider is that all these apparently descriptive words 'behave as if they encoded pro-concepts' because that is what they do in fact encode, that is, they fall in with the general class of indexicals (which are subject to a pragmatic process of saturation, of finding the appropriate semantic value in the context). This is not an attractive solution, for several reasons.

Indexicals constitute a specific small set of words, whose context-sensitivity is entirely systematic, while practically the whole of the descriptive vocabulary is modulated and in a non-systematic way. Indexicals come, in effect, with a slot and an instruction on the kind of thing to plug into the slot and the pragmatic process of slot-plugging is obligatory. The cases of context-dependence we are considering here are quite different: the pragmatically-derived concepts for a particular word can differ from each other in arbitrarily many ways and, even supposing we could set out the full range of parameters of variance, it would not be obligatory (or possible) to provide all of them with a semantic value on every occasion of use (Recanati 2004, Bosch 2007).

A second option is that a word type meaning is a ‘formal’ linguistic entity of some sort. For instance, in a discussion of the lexical semantics for classes of verbs, Glanzberg (2011) argues for a monadic conceptual root in a structural frame, along the following lines:

- (5) a. ‘X open Y’: [[X *act*] *cause* [*become* [Y <open>]]]
 b. ‘X hit Y’: [X *act*_{<hit>} Y]

It is important to note here that the structural frames for particular classes of verbs (‘open’ and ‘hit’ belonging to distinct classes) are grammatically determined and the components of the frames, *act*, *cause*, *become*, etc., are linguistic/grammatical elements, which are not identical to the ordinary concepts ACT, CAUSE, BECOME, etc. Whether the conceptual roots, ‘<open>’, ‘<hit>’, etc., are semantic or syntactic elements is left open (Glanzberg 2011: 9). This is clearly, then, a case of ‘wrong format’, that is, of linguistic semantic representations being a different kind of entity from conceptual (thought) representations, and the issue it raises, as for any other manifestation of this difference, concerns how we make the move from the one to the other in communication.

A third option is that words encode something schematic: a template for concept construction, a set of constraints, a rule for use, a sense-general meaning (or ‘archi-sememe’ or ‘super-concept’), as variously discussed by Ruhl (1989), Moravcsik (1994), and Atlas (1989). Ruhl, for instance, says: “... lexical meaning must be highly abstract (though still specific to a particular language), and thus highly formal, ... remote from all ambient contingencies” (1989: ix). It may be that Bosch’s (2007) ‘lexical concepts’, which do not determine an expressible content, also fall in here. No doubt, I am grouping together a disparate range of quite different manifestations of the general idea of stable word ‘meanings’ that underspecify but constrain the kinds of contents we can communicate when we use them. I cannot explore these different notions here, but note merely that something along these lines seems to be quite widely favoured.¹⁰

Under the ‘wrong format’ position, Recanati (2004) talks, on the one hand, of meanings that are too abstract and schematic, and, on the other, of

¹⁰Another account within this broad category, one that warrants detailed discussion, is that of Ruth Kempson, Eleni Gregoromichelaki and Christine Howes (2011), who characterise 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 the construction of representations of content by the interpreter in the process of utterance comprehension.

meanings that are too rich, incorporating a host of ‘semantic features’, which have to be whittled down (cancelled) by contextual/pragmatic considerations on any occasion of use in order to recover the particular sense expressed/communicated. Cohen (1993) takes a view of this latter sort, so that, for instance, the noun ‘rose’ has the features [plant, flower, petals, thorns on stem, beautiful, fragrant, highly valued ...] which are ordered in terms of centrality or prominence. When, for instance, the word is used metaphorically, or otherwise loosely, some of these features are cancelled. Evidently, Cohen includes in his set of ‘semantic features’ some components of what seem to be contingent encyclopaedic information (about roses, for instance), in an attempt to provide all the ingredients needed for different uses. Given that we have a vast store of encyclopaedic information about real world entities (e.g. roses) and that innovative uses of words (expressing genuinely new ad hoc concepts) are always possible, it seems that, in its bid to be comprehensive, this approach must ultimately collapse into a variety of what Recanati calls the ‘meaning eliminativist’ position.

A clearly eliminativist position is the recent ‘grab-bag’ model of Rayo:

With each expression of the basic lexicon, the subject associates a ‘grab-bag’ of mental items: memories, mental images, pieces of encyclopaedic information, pieces of anecdotal information, mental maps and so forth. With the expression ‘blue’, for example, a subject might associate two or three particular shades of blue, the information that the sky is blue, the information that my bicycle is blue, a memory of a blue sweater, and so forth. Different speakers might associate different grab-bags with the same lexical item. (Rayo forthcoming)

This looks very much like the kind of information associated with a word that is given as an encyclopaedic entry or material in a mental file in accounts that assume words encode concepts, as in relevance theory or Fodor (2008). It surely doesn’t qualify as the *linguistic* meaning or semantics of a word type – it is totally non-linguistic and largely contingent.¹¹

An appealing aspect of this approach is that it seems to provide an immediate and simple solution to the polysemy/metonymy problem. Referring to his example of the polysemous verb-noun ‘stop’, mentioned above, Rayo says: “One can place a few key items in one’s grab-bag for ‘stop’ – for instance, representations that bring to mind interfering, preventing, obstructing, closing – and let common sense and sensitivity to context take care of the rest” (Rayo forthcoming). And there is no need for different grab-bags for different grammatical categories: “a mental image that evokes obstruction, for example, can be used to render salient the action of closing a valve when interpreting ... ‘she stopped the flow of oxygen’ and to render salient a knob on a pipe organ when interpreting ... ‘she moved the stops to control the air flow into her organ’” (Rayo forthcoming). It is not too difficult to envisage a grab-bag for the word ‘novel’, which would include information about the stories, plots and characters

¹¹The instantiation of ‘meaning eliminativism’ that Recanati (2004: 146-151) sets out is rather different from the one discussed here. On the account he outlines, what a word form brings to the interpretation process is a stored collection of its previous uses and its interpretations in context.

that authors imagine, the written (or virtual) texts they may produce as a result, the publication, printing, selling and distribution processes, and the resulting physical copies. On different occasions of use, different selections are made from the grab-bag, in accordance with 'common sense and sensitivity to context', which I take to be a (somewhat cavalier) reference to the cognitive interpretive processes that a pragmatic theory seeks to explain.

An approach along these wholly pragmatic lines would put an end to the need for an array of semantic rules, even supposing they are formulable, and the futile attempts to decide which of the various senses of polysemous words to take as the basic one. Of course, it remains to be spelled out in detail exactly how the grab-bag selection process works, particularly how it can result in a concept with a truth-conditional content, but the basic intuition seems to be very much in keeping with the ad hoc concept construction process in relevance theory, as outlined in Section 2.

5. Conclusion

The goal of this paper was modest: to present a range of reasons for taking seriously the idea that words (or lexical forms) may not encode concepts or map directly to contentful entities, but rather come with meaning-relevant components that are different in kind from semantic values, that are intrinsically underspecified with regard to content, where a content is what is expressed/communicated by an individual's use of a word and so is only determinable on an occasion of use. Thus, this hypothesis is only worth exploring when coupled with a well-developed pragmatic theory that seems capable of providing a detailed account of how the concepts a speaker intends to express can be recovered by her addressee on the basis of such underspecified meanings or encyclopaedic information.

Abandoning a conceptual lexical semantics raises a host of new questions. First, there is the issue of maintaining a distinction between genuinely indexical words and these cases of word meaning modulation; on a non-conceptual construal of word type meaning, the pragmatic process of finding an appropriate semantic value is no longer optional, so cannot be distinguished on those grounds from the obligatory process of indexical saturation. Second, the approach has consequences for (relevance-theoretic) pragmatics, in that we can no longer think in terms of the narrowing or broadening of denotations (or of concept adjustment) as there is no linguistically-specified denotation to narrow or broaden (and no concept to adjust). All concepts occurring in communicated thoughts (explicatures) are pragmatically inferred and merely constrained by an encoded lexical schema/template or an array of activated encyclopaedic information (a grab-bag). In that sense, all concepts expressed or communicated are 'ad hoc'. Third, there is a robust intuition that many words have a 'literal meaning', that there is a particular concept (or concepts) which, among the others that the word may be used to express, is somehow privileged or basic. Whether and, if so, how this intuition is to be respected on any of these non-conceptual approaches to word type meaning needs to be addressed. Finally, the most pressing question is how to account for the move from a non-conceptual, non-semantic entity to a conceptual, contentful one. The project ahead, as I see it,

is to investigate each of these questions, in conjunction with various instantiations of the wrong format or eliminativist views of word meaning, within the explanatory pragmatic account provided by the relevance-theoretic framework.

References

- Asher, Nicholas. 2011. *Lexical meaning in context: A web of words*. Cambridge: Cambridge University Press.
- Atlas, Jay. 1989. *Philosophy without ambiguity*. Oxford: Clarendon Press.
- Borer, Hagit. 2005. *In name only: Structuring sense*, vol. 1. Oxford: Oxford University Press.
- Bosch, Peter. 2007. Productivity, polysemy, and predicate indexicality. In Balder ten Cate & Henk Zeevat (eds.), *Proceedings of the sixth international Tbilisi symposium on language, logic & computation*, 58-71. Berlin: Springer.
- Carston, Robyn. 2002. *Thoughts and utterances: The pragmatics of explicit communication*. Oxford: Blackwell.
- Carston, Robyn. 2012. Relevance theory. In Gillian Russell & Delia Graff Fara (eds.), *Routledge companion to the philosophy of language*, 163-176. London: Routledge.
- Carston, Robyn & Catherine Wearing. 2011. Metaphor, hyperbole and simile: A pragmatic approach. *Language and Cognition* 3(2). 283-312.
- Cohen, Jonathan. 1993. The semantics of metaphor. In Andrew Ortony (ed.), *Metaphor and thought*, 2nd edn., 64-77. Cambridge: Cambridge University Press.
- Fodor, Jerry A. 1998. *Concepts. Where cognitive science went wrong*. Oxford: Clarendon Press.
- Fodor, Jerry A. 2001. Language, thought and compositionality. *Mind and Language* 16. 1-15.
- Fodor, Jerry A. 2008. *LOT 2: The Language of Thought Revisited*. Oxford: Clarendon Press.
- Fodor, Jerry A. and Ernest Lepore. 1998. The emptiness of the lexicon: Reflections on Pustejovsky. *Linguistic Inquiry* 29(2). 269-288.
- Glanzberg, Michael. 2011. Meaning, concepts and the lexicon. *Croatian Journal of Philosophy* 11. 3-31.
- Kempson, Ruth, Eleni Gregoromichelaki & Christine Howes. 2011. Introduction. In Ruth Kempson, Eleni Gregoromichelaki & Christine Howes (eds.), *The Dynamics of Lexical Interfaces*, 1-20. Stanford: CSLI Publications.
- Lossius Falkum, Ingrid. 2011. *The semantics and pragmatics of polysemy*. London: UCL PhD dissertation.
- Moravcsik, Julius. 1994. Is snow white? In Paul Humphreys (ed.), *Patrick Suppes: Scientific Philosopher*, 71-85. The Netherlands: Kluwer Academic Publishers.
- Murphy, Gregory. 2002. *The big book of concepts*. Cambridge, MA: MIT Press.
- Nunberg, Geoffrey. 1979. The non-uniqueness of semantic solutions: polysemy. *Linguistics and Philosophy* 3(2). 143-184.
- Pustejovsky, James. 1995. *The generative lexicon*. Cambridge, MA: MIT Press.

- Rayo, Agustín. Forthcoming. A plea for semantic localism. *Noûs*
- Recanati, François. 2004. *Literal meaning*. Cambridge: Cambridge University Press.
- Ruhl, Charles. 1989. *On monosemy: A study in linguistic semantics*. Albany: State University of New York Press.
- Sperber, Dan & Deirdre Wilson. 1986/1995. *Relevance: Communication and cognition*. Oxford: Blackwell.
- Sperber, Dan & Deirdre Wilson. 1998. The mapping between the mental and the public lexicon. In Peter Carruthers & Jill Boucher (eds.), *Language and thought*, 184-200. Cambridge: Cambridge University Press.
- Textor, Mark & Nicholas Allott. forthcoming. Lexical pragmatic adjustment and the nature of ad hoc concepts. *International Review of Pragmatics*.
- Wilson, Deirdre & Robyn Carston. 2007. A unitary approach to lexical pragmatics. In Noel Burton-Roberts (ed.), *Pragmatics*, 230-260. Basingstoke: Palgrave.