MAGDALENA SZTENCEL (NEWCASTLE)

FROM WORDS TO CONCEPTS*

Relevance Theory (e.g. Carston 2002) has always subscribed to the underdeterminacy thesis in that natural language sentences do not encode propositions/thoughts. Relevance Theory's distinction between LINGUISTIC and REAL semantics and, derivatively, between LEXICAL and AD HOC concepts, is compatible with the underdeterminacy thesis as long as it is maintained that linguistic/lexical semantics does not have truth-theoretic properties. However, there is an instability in Relevance Theory (Sperber & Wilson 1995, Carston 2002) as to what linguistic semantics actually consists in. For example, Carston (2010) explicitly allows that at least some lexical concepts are 'full-fledged', i.e. they have truth-theoretic properties. I argue that because of Relevance Theory's adherence to the traditional Chomskyan double-interface view of linguistic expressions as objects constituted by phonological and semantic properties (PHON + SEM), and consequently because of their attribution of truth-theoretic properties to lexical concepts, it has become unclear in what form Relevance Theory still holds the underdeterminacy thesis.

I introduce the Representational Hypothesis (e.g. Burton-Roberts to appear *b*), a conceptual programme which rejects Chomskyan double-interface view of linguistic expressions (PHON + SEM) as problematic and anyway unnecessary to account for linguistic 'sound with a meaning'. In attributing no semantic properties to linguistic expressions, the Representational Hypothesis offers a more radical and consistent underdeterminacy thesis. I present the Representational Hypothesis' semiotic, wholly inferential, account of the relation between words and concepts, and argue that the mechanics of the representational account can be implemented in terms of Hintzman's (e.g. 1986) multiple trace theory of memory.

INTRODUCTION

Since Aristotle, language has been seen as a system that links sound and meaning (Chomsky 2005: 10). The Chomskyan tradition reconstructs this idea by attributing both phonology and semantics to linguistic expressions. On this view, linguistic expressions are taken to be double-interface syntactic objects constituted by phonological and semantic properties (Burton-Roberts 2009).

On this traditional account, phonological and semantic properties are in a part-part (mereological) relation with respect to the syntactic object they together constitute. This traditional double-interface assumption motivates the quest for lexical semantics, i.e. context-independent meaning encoded by linguistic signs and decoded by the linguistic system. In the double-interface tradition, linguistic signs have meaning in virtue of encoding lexical semantic properties: the word *cat* means CAT because it encodes the concept CAT. This effectively equates having meaning with having semantic properties.

A consequence of this assumption is that to account for meaning in language we need to posit two kinds of semantics – that of words and, on the assumption that thoughts too have semantics, that of thoughts. In Relevance Theory (e.g. Carston 2002), this idea is developed in terms of the distinction between linguistic/lexical semantics and real semantics, respectively.

This paper argues that the notion of linguistic semantics generally and of lexical concepts specifically is (a) highly problematic and (b) unnecessary in explaining meaning in language.

First, I introduce Relevance Theory's[1] distinction between linguistic and real semantics and explain its importance for Relevance Theory's underdeterminacy thesis. This is followed by a discussion of Relevance Theory's instability about the truth-theoretic nature of linguistic semantics and the implication this instability has for the underdeterminacy thesis. I discuss two processes of lexical concept formation in Relevance Theory – schema abstraction and lexicalisation – and point to some more problems and inconsistencies within Relevance Theory. I argue that because of Relevance Theory's adherence to the traditional Chomskyan double-interface view of linguistic expressions as objects constituted by phonological and semantic properties (PHON + SEM), and consequently because of their attribution of truth-theoretic properties to lexical concepts, it has become unclear in what form Relevance Theory still holds the underdeterminacy thesis. I then introduce the Representational Hypothesis (e.g. Burton-Roberts & Poole 2006, Burton-Roberts to appear *b*), a conceptual programme which rejects Chomskyan double-interface view of linguistic expressions (PHON + SEM) as problematic and anyway unnecessary to account for linguistic 'sound with a meaning'. In attributing no semantic properties to linguistic expressions, the Representational Hypothesis offers a more radical and consistent underdeterminacy thesis. I present the Representational Hypothesis' semiotic, wholly inferential, account of the relation between words and concepts, and argue that the mechanics of the representational account can be implemented in terms of Hintzman's (e.g. 1986) multiple trace theory of memory.

---

[1] Because Relevance Theory has received wide coverage in the literature, some basic knowledge of it is assumed.

LINGUISTIC SEMANTICS VERSUS REAL SEMANTICS

For Relevance Theory (e.g. Carston 2002, henceforth RT), the linguistic system 'maps' a phonetic representation into semantic representation, which is 'a structured (presumably conceptual) mental representation' (Carston 2002: 9). This semantic representation, or 'encoded linguistic meaning', is to be understood in terms of 'some appropriate notion of logical form, computed by linguistic decoding system' (ibid.). One of Carston's (2002) aims is to work out the details of the relation between the assumed linguistic meaning decoded by the language faculty and occasion-specific interpretations of utterances (i.e. propositions, or thoughts, communicated by utterances). This concern of Carston's reflects RT's distinction between linguistic meaning and speaker meaning. LINGUISTIC MEANING, often called LEXICAL CONTENT at the word level, is context-independent and refers to 'relatively stable meanings in a linguistic system, meanings which are widely shared across a community of users of the system' (Carston 2002: 19-20). In contrast, SPEAKER MEANING encompasses aspects of meaning conveyed by a speaker on a particular occasion. Hence, speaker meaning is context-dependent. This distinction, Carston (2002: 11) explains, is to be understood in terms of two types of cognitive processes: DECODING linguistic meaning and INFERRING speaker meaning.

As far as the linguistic input to pragmatic processing is concerned, Carston (2002: 28) explains that the logical form – or semantic representation of – an utterance is composed (partly[2]) by lexically encoded concepts, or LEXICAL CONCEPTS. RT represents lexical concepts by capitalised words; for example, the lexical concept encoded by the word *tired* is represented as TIRED. In utterance interpretation, the linguistic system first decodes such lexical concepts, from which AD HOC CONCEPTS are pragmatically inferred. This pragmatic inference is constrained by the principle of relevance (as defined by Sperber & Wilson 1995). *Ad hoc* concepts, Carston (ibid.) writes, are non-lexicalised concepts which are constituents of the proposition/thought expressed by the speaker. By 'non-lexicalised' Carston means not encoded as linguistic meaning. Consider (1) below.

(1) It's raw.

In broad terms, the linguistic meaning of (1) is that something is RAW, i.e. not cooked. From such linguistically en/decoded input, the pragmatics module infers the referent of *it*, for example the steak in front of the speaker, and it adjusts the meaning of *raw* so that it applies to things that have been cooked, but not for long enough. This adjusted meaning, i.e. the *ad hoc* concept communicated in a particular context, is represented as RAW*. The difference between

---

[2] In RT not all expressions encode lexical concepts, some encode procedural meaning.

RAW and RAW* shows, Carston (ibid.) argues, that speaker meaning cannot be determined by linguistic decoding alone; it has to be derived by pragmatic inference. In fact, for RT linguistic meaning UNDERDETERMINES utterance meaning in that it does not encode any proposition/thought at all. RT's linguistic underdeterminacy pertains to all levels: 'While sentences encode thought/proposition templates, words encode concept templates; it's linguistic underdeterminacy all the way down' (Carston 2002: 360).

Carston (2002: 58) explains that the inference from RAW to RAW* can be understood in terms of translating from one representational system to another. The difference between these representational systems is the following. By definition, the content of lexical concepts does not have truth-theoretic properties, i.e. it is not true or false of the world (it is non-propositional). In contrast, *ad hoc* concepts have truth-theoretic properties, i.e. they are related to objects and states of affairs in the mind-external world (they are propositional). In this sense the semantic content of *ad hoc* concepts is REAL as opposed to merely LINGUISTIC. Real semantics is identified by Carston (2002) as Fodor's (1998, 2008) Mentalese (i.e. the Language of Thought) – its elements are purely referential. Linguistic semantics, Carston (2002) maintains, is an incomplete, schematic Mentalese form. Being schematic and context-independent, it is not referential.

If it could be maintained, such a neat distinction between linguistic semantics and real semantics in terms of truth-theoretic value would be compatible with RT's underdeterminacy thesis; if linguistic meaning does not have truth-theoretic properties, it underdetermines the proposition expressed. However, as I will show, this radical understanding of linguistic underdeterminacy – in terms of truth-theoretic value or lack of it – is consistently undermined in RT by pervasive qualification. I turn to this now.

THE UNSTABLE NATURE OF LINGUISTIC SEMANTICS

In the spirit of RT's underdeterminacy thesis, Carston (e.g. 2002: 360) entertains the idea that words encode 'concept schemas' or 'pointers' to a conceptual space. As such, they are 'templates' for the construction of a fully propositional conceptual structure. On this proposal, concept schemas encoded by linguistic expressions are contrasted with the 'full-fledged' (truth-theoretic) nature of the semantics of thought/ real semantics.

In fact, the idea that thought is the only locus of real semantic properties has been one of the most important and radical, but not unproblematic, tenets of RT. RT's underdeterminacy thesis, and hence the overwhelming importance of pragmatic inference in grasping explicitly communicated meaning, rests on the assumption that en/decoded logical form is non-propositional. However, RT's

(Carston 2002) position is that at least some lexical concepts do in fact 'inherit' real semantic properties. According to Carston (2002: 362), in the case of natural kind terms like *cat*, *dog* or *lion*, there is a strong intuition that the element linguistically encoded is not merely a schema but a full-fledged concept.

The division of lexical meanings into schemas and full-fledged concepts brings with it a necessary distinction between two different kinds of semantics for these two kinds of lexical elements.[3] If natural kind concepts are full-fledged concepts and can therefore feature in thought, they must, for Carston, have referential semantics and thus, like other full-fledged concepts, must have truth-theoretic content. If natural kind concepts have referential/truth-theoretic content, they necessarily import this content into the level of sentences in which they occur. This undermines RT's underdeterminacy thesis which states that sentences are not truth-theoretic objects. In fact, the division of lexical content into schematic and full-fledged is not the only symptom of RT's instability regarding the nature of lexical content. The problem goes deeper.

Burton-Roberts (2005: 394) argues that RT's conceptual/propositional schemas (conceptual/propositional 'templates') are actually GENERAL concepts/propositions. Burton-Roberts' idea seems to be consistent with the fact that Carston often qualifies the claim that decoded logical form 'never' expresses truth-evaluable proposition by 'seldom if ever' or 'virtually never'. By doing this, Carston seems to allow that semantically general sentences like *Humans are mammals* do in fact encode propositions.

The problem is also manifest in RT's notion of explicature. Burton-Roberts (2005: 397) argues that RT's definition of explicature demands that truth-theoretic properties are encoded in natural language expressions. In RT, a proposition is an explicature (i.e. it is explicitly communicated) iff 'it is a development of (a) a linguistically encoded logical form of the utterance, or of (b) a sentential subpart of a logical form' (Carston 2002: 124). The problem with this definition is the following. The notion of 'development', on which the notion of 'explicature' rests, seems to be defined in terms of entailment such that X is a development of Y if X entails Y. But this means that logical form can, indeed must, be entailed. Since the relation of entailment – in fact any logical relation – can only hold between two propositional forms, this is inconsistent with RT's claim that logical form is non-propositional (Burton-Roberts 2005: 398).

Burton-Roberts (ibid.) notes that Sperber and Wilson (1995) too allow that decoded logical form enters into logical relations of contradiction and therefore implication. If that is the case, Burton-Roberts (2005: 396) argues, the output of linguistic decoding must be at least partly propositional. Again, if logical form is even partly propositional, the distinction of semantics into 'linguistic' and 'real' is seriously undermined.

---

[3] I am putting aside procedural meaning.

In fact, the instability as to whether decoded logical forms are propositional seems to have now become a defining feature of RT. This is clearly reflected in Carston (2010). Carston (2010: 242) maintains that the pragmatically inferred *ad hoc* concept: 'may be more specific or more general than the encoded concept; that is, its denotation may be either a proper subset or a superset of the denotation of the linguistically encoded concept […]'.

The idea of linguistic encodings as concepts, which themselves have denotations and which sometimes are more specific than *ad hoc* concepts pragmatically derived on their basis marks a shift away from RT's claim that encoded meaning does not have truth-theoretic properties. The point is that if concepts can be compared for specificity, they must have the same kind of semantics. It also suggests that vagueness versus specificity does not determine whether a concept counts as lexical or *ad hoc*. This strengthens Burton-Roberts' argument that in RT concept schemas are actually general concepts. And even though Carston (2010: 231) occasionally reminds us that 'by and large, sentences per se do not encode propositions', this has become an unsubstantiated claim.

Carston (2010: 244) is particularly revealing when she asserts that despite their heterogeneous nature (schemas, full-fledged concepts), all encoded meanings are atomic. She follows Fodor (1998, 2008) in assuming that conceptual correlates of words are not compositionally constituted; the meaning of a simple word like *cat* just is CAT. Since Fodor's (1998, 2008) referentiality with respect to conceptual content follows directly from this atomistic view of conceptual content,[4] we may assume that the content of such lexical concepts is referential for Carston too. Indeed, Carston (2010: 245) makes an unambiguous statement when she writes that lexical concepts are 'basic element[s] of the language of thought'.

With the linguistic-real division undermined, the distinction between lexical and *ad hoc* concepts is not any longer a difference between kinds of semantics. So what is it? In the next section, I look at the processes of SCHEMA ABSTRACTION and LEXICALISATION which in RT lead to the formation of two kinds of lexical concepts – concept schemas and full-fledged lexical concepts. I ask what it is about these processes that makes the resulting concepts distinct from *ad hoc* concepts. I look at the function that lexical concepts have in utterance interpretation.

SCHEMA ABSTRACTION AND LEXICALISATION

As Carston (2002: 365) puts it, linguistic meaning provides 'evidence (often rich and detailed evidence, but never a complete encoding, never a proof) of the thoughts being communicated'. Wedgewood (2007: 666) phrases it in terms of

---

[4] I will return to this issue in more detail.

shared content: 'RT notion of encoded meaning […] provides the logically necessary level of some infallibly shared content'.

There are two ways in which linguistic semantics is supposed to be shared in RT. The first pertains to context-independence – encoded meaning is shared across contexts. I refer to it as cross-context shareability. The second refers to shareability among speakers – in RT linguistic meaning is (widely) shared amongst members of the same speech community. I refer to it as cross-speaker shareability. The theoretical consequence of cross-context shareability is that the shared linguistic semantic content is automatically/deterministically decoded by the linguistic system. In RT, linguistic decoding is the first and necessary step in utterance interpretation. The theoretical weight of cross-speaker shareability falls on the assumption that it guarantees linguistic 'evidence' for the intended meaning, as quoted above. In what follows, I look at the significance of the two notions of shareability and their theoretical implications in RT with respect to the processes of schema abstraction and lexicalisation.

CONCEPT SCHEMAS: THEIR FUNCTION AND ACQUISITION

Recanati (1998: 630) makes reference to Hintzman's (1986) multiple trace theory of memory, which predicts that 'Words, as expression-types, do not have "meanings", over and above the collection of token-experiences they are associated with. The only meaning which words have is that which emerges in context'. Carston (2002: 375) makes reference to Hintzman's radical contextualism too, only to reject it in favour of a conservative view 'on which words do encode something, albeit something very schematic, which simply sends the system off to a particular region in long-term memory'. The schematic encoding that Carston talks about is, as discussed above, one proposed type of lexical meaning. As such it is necessarily accessed/ deterministically decoded in utterance interpretation. But there are problems.

Firstly, there are fundamental questions about the cognitive function of abstract/schematic lexical concepts and how they are acquired. Cross-context shareability demands that lexical semantics is a context-invariant, (relatively) stable 'meaning' of a word, which underlies all uses of a particular word. For example, for a word like *break*, lexical semantics is supposed to underlie such uses as *break a nose=damage, break the code=understand it and make it useless, break the strike=end, break the law=do something illegal, break the fast=start eating again*, *break someone's will=make someone lose control*, *break the news=tell the public,* etc. Even from this example alone we can see that for lexical semantics to underlie all uses of a given word it has to be extremely abstract or schematic, as in RT. The abstract nature of lexical semantics is actually the first

problem that we encounter: it is so abstract that it is almost never employed in thinking (e.g. Carston 2002: 363). It does not seem to have any function in our mental lives other than the purported mediation between a linguistic form and an *ad hoc* concept.

As mentioned, the abstract/schematic nature of lexically encoded meaning is compatible with RT's argument that logical form (along with its constituents) does not have truth-theoretic properties. But even Carston (2002: 363) herself admits that it is difficult to see how we ever manage to acquire such abstract schematic entities if they do not feature in our thinking about the world. Carston does not answer this question but simply assumes the following:

> There must be some process of abstraction, or extraction, from the particular concepts associated with the phonological form /open/ to the more general 'meaning', which then functions as a gateway both to the existing concepts of opening and to the materials needed to make new OPEN* concepts which may arise in the understanding of subsequent utterances. (Carston 2002: 364)

However, as noted by Hintzman (1986) and Burton-Roberts (2007), there is a problem with the acquisition of lexical schemas. Lexical semantics is supposed to be acquired in experience and, in Carston's terms, it provides a 'gateway' to the understanding of the meaning of a word in context. It guides the hearer to the intended context-dependent meaning of a word. But lexical semantics is abstracted from, and hence presupposes prior understanding of utterances. In other words, the acquisition of lexical semantics is *post hoc*. But if the acquisition of lexical semantics is *post hoc* and presupposes prior understanding of utterances, it cannot be a necessary step in utterance interpretation.

In a nutshell, the problem is that positing concept schemas as lexical semantics amounts to positing constructs which are either (a) impossible to acquire, since it is impossible to employ them in thinking or (b) possible to acquire, but not necessary in understanding an utterance. If (a), concept schemas cannot function as cross-context shared content. If (b), their *post hoc* acquisition cancels the condition that they be necessarily (i.e. deterministically) accessed.

I will later argue that conceptual schemas – even *when* they exist – are not linguistically decoded, but pragmatically inferred. In the meantime, I look at the process of lexicalisation.

FULL-FLEDGED LEXICAL CONCEPTS AND FODOR'S NATURALISTIC LEGACY

As discussed, both schematic and full-fledged lexical concepts are, for Carston, atomic. At the same time, however, she maintains that such 'atomic' concepts may make available three kinds of information: (i) logical content, which

captures analytic relations; (ii) encyclopaedic or general knowledge; and (iii) lexical properties, which specify phonological and syntactic properties of the linguistic form which encodes a given lexical concept. By assumption, the information specified in (i), (ii) and (iii) is non-content constitutive, which means that the content of a concept is independent of it.[5]

Carston (e.g. 2010: 246) is ambivalent about the logical-encyclopaedic division and admits that the distinction between (i) and (ii) is 'rather controversial'. Leaving aside (iii), this allows us to group (i) and (ii) under the label 'non-content-constitutive information' and to contrast it with conceptual content. In RT, conceptual content is en/decoded by the linguistic system, whereas the information specified in (i) and (ii) is non-content-constitutive precisely because it is *not* en/decoded – by assumption, it has to be inferred (e.g. Carston 2002: 11). In drawing this division between conceptual content and non-content-constitutive information Carston follows Fodor (1998, 2008), who writes:

> […] concepts have both referents and a congeries of beliefs (etc.) in which they are embedded. It's just that, whereas the former has to do with the content of the concept, the latter has to do with its (e.g. inferential) role in mental processes. The distinction between these is independently motivated; *content is what composes*, and inferential roles and the like do not. (Fodor 2008: 87-88)

It is important to notice that Fodor's atomism about conceptual content and, relatedly, his adoption of referentialism are closely tied up with his quest for a *naturalistic*[6] account of conceptual content. Fodor's (e.g. 1998: 30-34) theory is supposed to provide an account of content which is shared *across the species*; Fodor wants an account of *the* human mind, not any particular mind.

Given his naturalistic pursuits, Fodor rejects (i) as content-constitutive because it is easily shown, he argues, that analytic relations are merely epistemic. Fodor (1998: 74) points out that even though number two, for example, seems to be necessarily a prime number, it is undeniably possible for someone to have the concept TWO but not the concept PRIME. This shows that the inference from *two* to *prime* is not actually necessary, i.e. that PRIME cannot be a necessary component of TWO. The relation between the concept PRIME and the concept TWO is epistemic, hence *non*-naturalistic and hence *not* semantic. Fodor thus rejects the analytic-synthetic distinction, and groups (i) and (ii) as non-content-

---

[5] Groefsema (2007) observes that atomicity is incompatible with the notion of concept 'modulation' or 'adjustment'. Based on this and many other insightful observations, she defends a view that conceptual content should be seen as constituted by logical and encyclopaedic entries. Due to space restrictions I will not discuss Groefsema's paper here, relevant though her arguments are to the position I am defending here.

[6] In Fodor's (e.g. 1998: 28, 34) terms, a theory of concepts should meet the publicity constraint on concepts, i.e. a constraint by which a theory should predict that we all share types of the same primitive concepts. In order to satisfy this publicity constraint, Fodor (e.g. 1998: 148-150) 'naturalises' reference in terms of his mind-dependence thesis.

constitutive, because – being the locus of individualistic and thus holistic be-
liefs – (i) and (ii) cannot give him a naturalistic account of conceptual content.
Ultimately, it is Fodor's search for a naturalistic (i.e. species-shared) conceptual
content that leads him to reject the analytic-synthetic distinction, adopt atomism
and, consequently, a referential account of content.

The problem for RT is that Fodor's species-shared notion of conceptual con-
tent is incompatible with RT's standard, social-externalist assumption that lexi-
cal content is shared amongst the members of the same speech community (e.g.
Sperber & Wilson 1998). Fodor's notion of content is also incompatible with
RT's underdeterminacy thesis precisely because such content has truth-theoretic
(referential) properties. It seems thus that RT is following Fodor's conclusion
about conceptual content without agreeing on the premise. However, despite the
claims to adopt Fodor's theory, Carston (2010) diverges from it in a substantial
way. This is where the process of lexicalisation enters the stage.

Carston (2010: 250) suggests that the most plausible assumption is that both
lexical and *ad hoc* concepts are atomic elements of the Language of Thought
(henceforth LOT).[7] The difference between them is attributed to the process of
'lexicalisation' of certain concepts – presumably, a process by which some *ad
hoc* concepts are elevated to the status of lexical content. Full-fledged lexical
concepts are thus *ad hoc* concepts which have been 'progressively conventiona-
lised', lost their *ad hoc* status and become lexicalised, i.e. encoded as linguistic
meaning (Carston 2010: 244).

For Carston, an atomic lexical concept is a 'conceptual address' or a 'file
name' to which some individualistic/holistic information is, by assumption,
non-constitutively attached. *Ad hoc* concepts, Carston (2010: 247) argues, are
contextually derived from such an individualistic/holistic pool of information
attached to lexical concepts – a step constrained by the principle of relevance.

Now, Carston (e.g. 2002: 214, fn 31) emphasises the following difference
between Fodor's account and RT. Fodor, she argues, endorses a code model of
communication whereby the concepts communicated are the same as the concepts
en/decoded. On RT, however, there is only a partial mapping between words and
concepts since large part of the conceptual repertoire is not lexicalised.

This, in effect, is a difference in conceptual architectures. For Fodor, LOT
hosts concepts whose content is species-shared. For him, all associated informa-
tion about perceived logical and encyclopaedic relations is not content-consti-
tutive, but idiosyncratic, holistic and thus irrelevant for semantics. Semantics,
for Fodor, is free from individualism/holism (2008: 88). For Carston, however,
concepts are more numerous than those identified as word meanings. This is be-
cause contextually restricted accessing of the logical/encyclopaedic information

---

[7] Carston (2010) does not explicitly say that schemas are concepts, but, based on the arguments
presented in the preceding sections, I assume it follows from her account.

results in constructing *ad hoc* concepts. In this sense, Carston's LOT (unlike Fodor's LOT) is populated by pragmatically (i.e. holistically) derived atomic concepts. Carston's conceptual content (unlike Fodor's content) is *not* free from individualism/holism. Carston thus substantially departs from Fodor's notion of conceptual content.

The conceptual architecture proposed by Carston (2010) clearly indicates that she does not seek semantic content which is species-shared. Since *ad hoc* concepts are pragmatically derived on the basis of individualistic/holistic beliefs and if it is (some of) such holistically derived concepts that are lexicalised, i.e. turned into lexical concepts, they cannot be shared amongst the species. Unfortunately, however, Carston is inconsistent on this. Despite appearing to reject Fodor's position on shareability of content, Carston still maintains that the constructed *ad hoc* concepts – and, derivatively, full-fledged lexical concepts – have referential semantics. Since Fodor's 'atomic and referential' is inter-dependent with 'species-shared', Carston's account of conceptual content is contradictory. The assumed referential nature of lexical concepts is also incompatible with RT's underdeterminacy thesis. In essence, RT's instability about linguistic/lexical semantics can be summarised as follows:

> By 'linguistic semantics', I suggest, RT attributes either too little or too much in the way of semantics to particular languages. It is too little to be consistent with traditional assumption that particular languages have semantics as ordinarily understood, and too much to be consistent with LOT being the sole locus of real semantic properties. (Burton-Roberts 2007: 91)

The argument advocated here is that RT attributes too much semantics to words. No matter what kind of content RT attributes to words, there always is a problem for RT. The first assumption of RT – that lexical meanings are not concepts, but rather some schematic objects with no truth-theoretic properties – is compatible with RT's underdeterminacy thesis but difficult to maintain. In line with Burton-Roberts (2005), I have argued that even in RT's terms conceptual schemas are actually general concepts. The second assumption of RT – that lexical meanings, whether abstract or full-fledged, just are concepts – is incompatible with RT's underdeterminacy thesis and, furthermore, involves a contradiction.

The more positive upshot of this situation is that RT seems to be withdrawing the claim that there is non-truth-theoretic linguistic semantics. This effectively amounts to positing just one kind of semantics – that is, real semantics. In what follows I argue that positing just one kind of semantics – real semantics – can, in fact, be compatible with a more radical underdeterminacy thesis. The issue is not any longer what sort of semantics lexical concepts have, but whether positing lexical concepts and the process of deterministic decoding is even necessary to account for meaning in language. In the next section, I develop this idea within the framework of the Representational Hypothesis.

THE REPRESENTATIONAL HYPOTHESIS

Any theory of meaning is supposed to explain mutual understanding between interlocutors against the background of holistic, idiosyncratic beliefs. This is a huge task and lexical semantics is supposed to aid it by providing a (relatively) stable mutual core of meaning shared by interlocutors in communication. At the cost of undermining the underdeterminacy thesis, RT's notion of lexical content is designed to offer such a piece of 'linguistic evidence'. Since in RT (especially Carston 2010) this linguistic evidence is constituted by truth-theoretic content, it has become difficult to pinpoint the way in which encoded meaning is supposed to underdetermine the proposition expressed. The encoded and inferred kinds of concepts do not differ in terms of truth-theoretic value – the difference lies solely in the assumed difference between the deterministic process of decoding and context-constrained process of inferring.

As mentioned in the Introduction, RT's distinction between deterministically decoded lexical concepts and *ad hoc* concepts is motivated by the traditional double-interface assumption whereby linguistic signs have meaning in virtue of being partly constituted by semantic properties. In RT it is in virtue of the encoding-inferring distinction that linguistic signs purportedly have some specifically linguistic meaning. In this sense, the encoding-inferring distinction is dictated by the double-interface view.

However, Burton-Roberts (2007, 2009, to appear *a*, to appear *b*; with Poole 2004, 2006) argues that the double-interface idea is conceptually impossible, unexplanatory and unnecessary.[8] One of the problems is that even on Chomskyan Minimalist terms words qua double-interface objects (SEM + PHON) cannot actually exist. The Minimalist principle of Full Interpretation states that semantic properties are interpreted only at LF, and phonological properties are interpreted only at PF. Because words as double-interface objects are not interpretable at any of the two interfaces, the two kinds of properties are separated at Spell-Out to satisfy the principle of Full Interpretation. Since LF and PF are the only levels of representation, this means that Chomskyan words (SEM + PHON) and objects composed of them do not actually exist. The question that Burton-Roberts (2009, to appear *a*) asks is why posit such double-interface objects if it is acknowledged that the two properties are sortally incompatible and hence mutually uninterpretable.

Furthermore, the double-interface idea is conceptually unexplanatory because positing the part-part (mereological) relation between phonological and semantic properties is insufficient to account for meaning. In order to account for meaning in language, the mereological relation has to be supplemented (as

---

[8] I will not rehearse Burton-Roberts' arguments here, merely summarise the main ideas. The interested reader is directed to the sources quoted.

it is for Saussure) by an arbitrary, i.e. non-natural but conventional, *semiotic* relation, whereby the sound properties function as a *signifier* and the meaning properties (concepts) function as a *signified*. Burton-Roberts argues that the explanatory power of the semiotic relation renders the mereological relation conceptually unnecessary. In what follows I concentrate on this, perhaps the most important, criticism of the double-interface idea – the claim that it is unnecessary. I introduce the Representational Hypothesis, a framework which is motivated by the problems with the double-interface tradition and which develops the semiotic idea.

The claim of the Representational Hypothesis (henceforth RH) is that utterances of linguistic expressions – i.e. sounds utilised by speakers in communicating thoughts – are symbolic signs. Like other signs, they involve communicative intention as well as its recognition and convention. Like other symbolic[9] signs, they are meaningful for someone when the communicative semiotic intention is recognised.

The RH's semiotic account of the relation between sounds and conceptual structures brings with it a radical change in the understanding of what it is for a word to mean something. In the RH, having meaning does not equal having semantic properties. Words (i.e. sounds) clearly have meaning but do not have semantic properties as a constitutive property; words themselves have no conceptual attributes. Only thought has conceptual/semantic properties.

The RH's claim that thought is the only locus of semantics is closely tied up with its rejection of the traditional Saussurean-Chomskyan view of a linguistic sign. In rejecting the double-interface (SEM + PHON) view of linguistic expressions, the RH adopts the semiotic account of C. S. Peirce, where a sign is *other* than what it signifies (Burton-Roberts 2007, to appear *b*). On the RH, sounds of particular languages function as symbolic signs which are used to conventionally REPRESENT[10] internalist conceptual/semantic content.

Burton-Roberts' argument can be illustrated by the following example. A 'no stopping' road sign (i.e. a blue circle with a red cross) clearly has no semantic/conceptual properties. Nevertheless, when intentionally used it is meaningful to someone because it gives rise to the conceptually/semantically constituted thought that one should not stop the car in a particular area. The meaning of this sign, in fact of any sign, lies in the relation it has to a thought in a cogniser's mental world.

In a similar way, words, in virtue of being symbolic signs, do not have semantic/conceptual properties. However, they are meaningful to us when intentionally used because they give rise to a particular conceptually/semantically

---

[9] 'Symbolic' is used in contrast to 'indexical', as in the work of C. S. Pierce.

[10] The term 'representation' here, unlike in generative literature, is used in a relational sense – 'representation' is 'representation of something else'; in the linguistic context, an acoustic event is used to represent, and hence a representation of, a conceptual structure.

constituted thought.[11] In this sense, the RH makes an unambiguous distinction between *meaning* and *semantics*.

In the RH *meaning* and *semantics* are different but related in the following way. Meaning is a semiotic RELATION which holds between X (e.g. a phonetic phenomenon) and what has semantic content Y (thought and only thought). So, phonetic phenomena have meaning *for someone* in virtue of being in an intended, conventional, semiotic, relation to semantic entities. Sound is articulated in aid of conventionally representing expressions generated by an internalist system which the RH identifies as LOT.

One of the RH's central claims is that natural languages are Conventional Systems for the Physical Representation (CSPRs) of LOT. Crucially, the RH (e.g. Burton-Roberts & Poole 2006) distinguishes between the object represented (*representatum*) and the physical phenomena that it is represented by (*representans*). *Representatum* (x) is what is generated by LOT, i.e. conceptual structures. CSPRs (i.e. particular languages) define what counts as a well formed *representans* ($R(x)$). The relation between x and $R(x)$ is that of conventional representation, where $R(x) \neq x$. The *representatum* (x) is innate and invariant across the species.[12] By contrast, how it is represented in particular languages (CSPRs) involves massive variation. It is precisely because the representational relation between $R(x)$ and x is not natural but conventional that different CSPRs are constituted by different representational conventions.
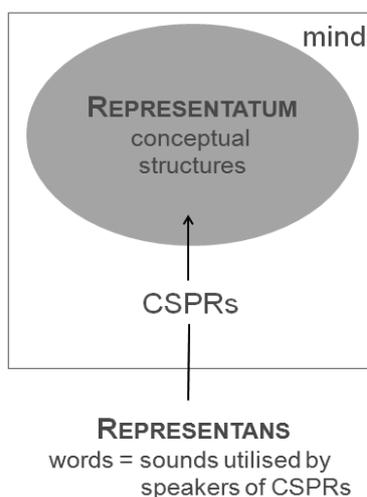

Figure 1. Representational architecture of mind

---

[11] I will shortly discuss the relevant issue of meaning potential.

[12] The issue of what is meant by 'innate and invariant' will be discussed in the next section.

On the RH (Burton-Roberts, to appear *b*), a concept represented in using a word, or concatenation of words, is taken to be (typically) structured.[13] It is helpful, I think, to identify the primitive components of such a conceptual structure as Jackendoff's (2002) 'quarks' – i.e. concepts 'ineffable' in isolation. Such conceptual quarks and structures defined over them (generated by LOT) are not acquired but innate. Burton-Roberts (to appear *a*, fn 16) argues that concepts are prior to and hence independent of language acquisition during which they are activated and accessed by proxy (i.e. labelled by a representational label[14]). Note that on the RH conceptual 'quarks' are posited on independent principled grounds: if sound is to function as a sign, there must be something it is a sign of, namely a signified, independently of the fact of signification. In other words, that there be a concept to signify is a precondition for signification.

The implications of the RH's claim that words do not have meaning as a constitutive property are very radical. The way in which CSPRs mediate between sounds and conceptual structures (Fig.1) does not involve positing double-interface objects and, consequently, it does not involve positing a distinction between deterministic decoding and contextually-constrained inferring. So how do CSPRs mediate between acoustic events and conceptual structures?

The RH's stand on the acquisition of access to conceptual structures in a particular language is compatible with Hintzman's (1986, 1988) multiple trace theory of memory. This model predicts that words do not encode lexical meaning/semantics. In the final section, I argue that the RH should adopt Hintzman's model not only for the acquisition of access to conceptual structures, but also for modelling utterance interpretation.

THE REPRESENTATIONAL HYPOTHESIS AND HINTZMAN'S MULTIPLE TRACES

The overarching question that Hintzman (1984, 1986, 1988, 2008) is concerned with is how abstract (generic) knowledge is related to specific (episodic) experience. On Hintzman's view, each and every experience, including linguistic experience, produces individual traces in long term memory. Such memory traces are constituted by configurations of primitive properties which are innately

---

[13] Due to space considerations, I leave out a relevant discussion of some problems with Fodor's atomicity and referentialism. My general point with respect to RT and conceptual content is that they should reject Fodor's atomicity and referentialism. Not only is adhering to atomicity and referentialism incompatible with their notion of content which is 'merely' shared by members of the same speech community (i.e. not species-shared), but, more fundamentally, it is highly problematic and circular in Fodor's own terms. The reader is directed to Sztencel (to appear) for further details.

[14] The labelling metaphor is consistent with Hintzman (1986: 412) and Chomsky (2000: 61).

specified. On this view, generic knowledge does not have a special status and is not stored in a functionally separate memory system, but can be retrieved *on-line* from a pool of memory traces.

Earlier I discussed concept schemas, i.e. abstract, very schematic concepts, as potential candidates for RT's lexical encodings. In RT, such schematic concepts are deterministically decoded rather than inferred. Hintzman, however, argues that abstract concepts are not retrieved by any linguistic default. Just like more specific concepts constituted by temporal or spatial information, abstract concepts can be retrieved from a pool of memory traces *only when the context requires it*. Hintzman's model has serious implications for linguistic theory as it dispenses with the encoding-inferring distinction. I explain this in the remainder of this section.

On Hintzman's theory, each linguistic experience is stored as a memory trace. As for the question of what memory traces consist in, Hintzman assumes that experiences are internally represented as an active configuration of primitive properties. By 'primitive properties' Hintzman means anything from modality-specific sensory features (e.g. basic colours and odours), simple emotional tones, properties accessible by more than one modality (e.g. intermittency, spatial location) to primitive abstract relations (e.g. before, same as). Similar experiences share certain configurations of properties. Hintzman (1986: 412) argues that such primitive properties are distinct from the ability to 'label' them in that they are not acquired through experience.

Hintzman (1986) distinguishes between *primary memory* (PM) and *secondary memory* (SM). PM is the active representation of the current experience, whereas SM is a pool of largely dormant memory traces. PM and SM interact in the following way. The active configuration of primitive properties in PM constitutes a RETRIEVAL CUE or PROBE which is sent to all traces in SM. The probe which is sent to SM activates traces in SM according to their similarity to the probe. PM then receives a single reply or ECHO from SM. The echo that emanates back from SM is a pattern of most strongly activated properties and it is the echo that, for Hintzman, constitutes the meaning of a word on a particular occasion of use. Depending on the structure of the probe, the information retrieved from SM will be of different degrees of abstractness.

In a linguistic context, Hintzman's model predicts that every communicative event to which a person attends – such as hearing a given word – will leave a new memory trace. Such a trace will co-exist in memory with other occurrences of the same word. Using the 'label' metaphor mentioned above, we can say that there will be aggregates of memory traces associated with, or attached to, acoustic labels.

Hintzman (1986, 2008) argues that the process of echo retrieval can retrieve 'the essence' of what, for example, dogs are from individual memory traces. When cued (e.g. when asked to think of a definition of a word, or on hearing

a generic statement), a generic concept of a dog can be produced on-line by cumulative activating of all traces and cancelling out the properties that are not shared by the traces. On this view a schema is 'a temporary, dynamic structure that springs into being when a retrieval cue occurs' (Hintzman 1986: 424). The crucial point is that the retrieval of such abstract/schematic echo – like all instances of echo retrieval – is necessarily context-dependent. In other words, echo retrieval can yield 'different nuances, different levels of abstraction, or entirely different meanings of a word by addressing different subsets of stored contextual features' (2008: 25). This means that on a generic use of a word the individuating properties of traces, such as temporal and spatial location properties, will be cancelled out. The experience, Hintzman (1984: 241) remarks, will be abstract and devoid of specific details. The crux of Hintzman's theory is that the retrieval of such an abstract/schematic concept does not, in any sense, happen by some default – the process underlying it is the same as the process underlying the retrieval of more specific concepts.

This context-sensitivity follows from the definition of the probe. The probe, Hintzman (1986: 420) emphasises, consists not only of a word (qua-acoustic-event) but also of the context of its use. The echo retrieved by such context-sensitive probes, i.e. a function of the particular subset of episodic traces activated by the probe, is thus necessarily context-sensitive too, whether schematic or not.

In Hintzman's model then there is no room for context-independent conceptual schemas necessarily mediating utterance understanding. The general mechanism of echo retrieval handles cases where conceptual schemas are retrieved and cases where they are not retrieved. On Hintzman's model, the existence of conceptual schemas (as memory traces) does not amount to the existence of lexical semantics understood in terms of a mental representation retrievable by a linguistic default of decoding.

The idea that both lexical schemas and more specific conceptual structures are retrieved by the very same contextually-constrained process is particularly important in the context of Relevance Theory. As mentioned, RT's distinction between linguistic and real semantics is cashed out in terms of two types of cognitive processes – deterministic decoding of linguistic meaning (whether schematic or full-fledged) and pragmatic inferring of speaker meaning. Hintzman's model effectively undermines this distinction, making the process of utterance interpretation radically contextualist and wholly inferential (this is compatible with arguments in Burton-Roberts, to appear *b*).

The upshot of combining the Representational Hypothesis with Hintzman's model is this. If the RH's *representatum* can be identified as Hintzman's echo, it becomes possible to explain how contextually available information constrains the search for speaker-intended meaning. Hintzman's aggregate level allows us to explain (a) the words' *meaning potential* without invoking the problematic and unnecessary notion of linguistic semantics/decoding and, therefore, (b) to constrain

the range of possible interpretations (echo/*representatum*).[15] This allows us to maintain that thought is the only locus of semantic properties without undermining the underdeterminacy thesis. Linguistic signs radically underdetermine thought because the process of getting from linguistic signs to thoughts is wholly inferential.

Importantly, the account presented here is supported by the consideration of compositionality, a principle which is supposed to explain productivity (i.e. the infinite expressive power) of language. In the context of RT's distinction between lexical and *ad hoc* concepts, the compositionality principle should operate at the linguistic semantic level if it is to explain productivity of language. Accordingly, RT's compositionality principle operates to combine the lexical semantic content of linguistic expressions into a logical form. The resulting structurally complex logical form delivered by the linguistic module serves as an input to pragmatic processing. However, there is considerable evidence which shows that compositionality does not take place at the context-independent level.

For example, Recanati (2005) argues that the compositionality principle applies after the pragmatic processes have done their work. In Recanati's words, pragmatic processes do not operate 'globally' on a compositionally constituted 'output of the grammar'. The order is actually reverse to what Relevance Theory predicts – pragmatic processes are at play before the compositionality process applies. Consider (2) and (3) below (taken from Recanati 2005):

(2) There's a lion in the courtyard.
(3) There's a stone lion in the courtyard.

As for (2), we understand that what is said to be in the courtyard is not a real animal but a representation, or statue, of a lion. The pragmatic inference that leads us to interpret the word *lion* in the sense of a statue is often referred to as REFERENCE TRANSFER. In the case of (3) too, we are dealing with reference transfer – we understand that what is said to be made of stone is a representation, or statue, of a lion, not a real animal. This fact, Recanati (ibid.) argues, shows that reference transfer must take place before the composition rule applies to the noun-noun (*stone lion*) construction. The evidence is this. The interpretation that we get for (3) is: *(a representation of a lion) that is made of stone*. But this interpretation is only possible if reference transfer occurs before the two expressions, *stone* and *lion*, are combined. If, however, reference transfer applied globally, i.e. after compositionality process was applied at the lexical semantic level, the interpretation we would get is: *a representation of (a lion that is made of stone)*. The absurdity of the result, Recanati argues, seriously undermines the view that compositionality applies at the lexical semantic level.

It may be argued that this criticism is not applicable to RT due to RT's underdeterminacy thesis and its emphasis on the pragmatic contribution to grasping

---

[15] This account is compatible with and, I believe, has to be supplemented by Bilgrami's (1992) locality of content thesis. Due to space restrictions, I do not discuss it here.

the proposition explicitly expressed. However, in the light of the discussed RT's instability about the propositional nature of logical form and truth-theoretic properties of lexical concepts, underdeterminacy thesis becomes an empty argument against Recanati's criticism. In my opinion, what Recanati's argument challenges is not only the issue of compositionality at the lexical semantic level, but, more fundamentally, the utility of the linguistic semantic level as such.

CONCLUSION

I have argued that the notion of linguistic/lexical semantics is highly problematic and conceptually unnecessary to account for meaning in language. An alternative, semiotic, wholly inferential, account within the framework of the Representational Hypothesis has been presented. I have argued that the details of the representational account should be cashed out in terms of Hintzman's (e.g. 1986) multiple trace theory of memory.

If Burton-Roberts and Hintzman are right, there is no interesting sense in which the notion of a word as a double-interface object could be reconstructed. If words are supposed to steer interpretation or direct hearers to conceptual space, then words, on such an account, can only be reconstructed as acoustic events. The way in which words-qua-acoustic-events *point* to specific regions of conceptual space is restricted by the contents of the trace aggregate level (which can be understood as the meaning potential) and context-specific considerations.[16] This, in my opinion, is a very promising way to go.

REFERENCES

BILGRAMI, A. (1992): *Belief and Meaning*, Cambridge, Massachusetts: Basil Blackwell Ltd.

BURTON-ROBERTS, N. (2005): Robyn Carston on semantics, pragmatics and 'encoding', *Journal of Linguistics* 41, 389-407.

BURTON-ROBERTS, N. (2007): Varieties of semantics and encoding: negation, narrowing/loosening and numerals, in: BURTON-ROBERTS, N. ed. *Pragmatics*. Basingstoke: Palgrave-Macmillan, 90-114.

BURTON-ROBERTS, N. (2009): The Grounding of Syntax – and More, *Newcastle Working Papers in Linguistics* 15, 21-39.

---

[16] Hintzman's operation of echo retrieval whereby primitive properties are activated according to their similarity to the probe can be seen as governed by RT's principle of relevance, which restricts utterance interpretation to that which yields a positive cognitive effect (a stimulus is relevant if it makes a difference in an individual's representation of the world) at the least processing effort (the required use of perception, memory and inference, in establishing the most relevant conclusion). No lexical semantics is needed.

Burton-Roberts, N. (to appear *a*): On the grounding of syntax and the role of phonology in human cognition, *Lingua*.

Burton-Roberts, N. (to appear *b*): Meaning, semantics and semiotics, in: Capone, A. (ed.), *Perspectives on Pragmatics and Philosophy*, Springer.

Burton-Roberts, N. / Poole, G. (2004): Syntax vs. phonology: a representational approach to stylistic fronting and verb-second in Icelandic, *Lingua* 116, 562-600.

Burton-Roberts, N. / Poole, G. (2006): "Virtual conceptual necessity", feature-dissociation and the Saussurian legacy in generative grammar, *Journal of Linguistics* 42, 575-628

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

Carston, R. (2010): Explicit communication and 'free' pragmatic enrichment, in: Soria, B. / Romero, E. (ed.), *Explicit communication: Robyn Carston's Pragmatics*, Basingstoke: Palgrave, 217-287.

Chomsky, N. (2000): *New Horizons in the Study of Language and Mind*, Cambridge: Cambridge University Press.

Chomsky, N. (2005): Three factors in Language Design, *Linguistic Inquiry* 36 (1), 1-22.

Fodor, J. (1998): *Concepts: Where Cognitive Science Went Wrong*, Oxford: Oxford University Press.

Fodor, J. (2008): *LOT 2: The Language of Thought Revisited*, Oxford: Oxford University Press.

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

Hintzman, D. (1984): Episodic versus semantic memory: A distinction whose time has come – and gone? *Behavioral and Brain Sciences* 7 (2), 240-241.

Hintzman, D. (1986): "Schema abstraction" in a multiple trace memory model, *Psychological Review* 93, 411-28.

Hintzman, D. (1988): Judgments of frequency and recognition memory in a multiple-trace memory model, *Psychological Review* 95, 528-551.

Hintzman, D. (2008): Memory from the outside, memory from the inside, in: Gluck M. A. / J. R. Anderson / S. M. Kosslyn (eds.), *Memory and Mind*, New York: Lawrence Erlbaum, 15-30.

Jackendoff, R. (2002): *Foundations of Language: Brain, Meaning, Grammar, Evolution*, Oxford: Oxford University Press.

Recanati, F. (1998): Pragmatics, *Routledge Encyclopedia of Philosophy. Vol 7*. London: Routledge, 620-633.

Recanati, F. (2005): Literalism and Contextualism: Some Varieties, in: Preyer, G. / Peter G. (eds.) *Contextualism in Philosophy*, Oxford: Oxford University Press.

Sperber, D. / Wilson, D. (1995): *Relevance: Communication and Cognition*, Oxford: Blackwell; Cambridge, MA: Harvard University Press.

Sperber, D. / Wilson, D. (1998): The mapping between the mental and the public lexicon, in: Carruthers P. / Boucher J. (eds.) *Language and Thought: Interdisciplinary Themes*, Cambridge: Cambridge University Press, 184-200.

Sztencel, M. (to appear): Against referential semantics, in: Frath P. et al. (eds.) *Res-per-nomen III: La référence, la conscience et le sujet énonciateur / Reference, Consciousness and the Speaking Subject*. Editions et Presses Universitaires de Reims.

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