

Varieties of semantics and encoding: negation, narrowing/loosening and numerals.¹

Noel Burton-Roberts

[To appear in N. Burton-Roberts (ed.) *Advances in Pragmatics*. Palgrave]

1. Varieties of Semantics?

Semantics is ‘conceptual-intentional’. It articulates the connection between a language and the world. As such, it is truth-theoretic, dealing with the truth-conditions of truth-evaluable entities. Thoughts are semantic in this sense. The language in which thoughts are couched is the ‘Language of Thought’ (LoT).

This much, I believe, is generally agreed upon if only because it skates over much we don’t understand. It is certainly part of the picture assumed in Relevance theory (RT). But RT makes two claims that, taken together, are difficult to fathom.

Claim 1 is that LoT is not just a locus, but the (only) locus, of semantic properties. Taken seriously, this implies that the particular languages we speak don’t have semantics. However, RT’s second claim at least obscures Claim 1. **Claim 2** is that particular languages do in fact have semantics. RT seeks to reconcile these claims by distinguishing two varieties of semantics, ‘real’ and ‘linguistic’. Particular languages, for RT, have only ‘linguistic’, not ‘real’, semantics.

While ‘real semantics’ is clear enough—it’s semantics as outlined above—it is less clear what ‘linguistic semantics’ is. The force of RT’s ‘linguistic’~‘real’ semantics distinction is essentially negative: to withhold ‘real’ (truth-theoretic) semantic properties from the expressions of particular languages. The oft-repeated claim of RT is that the sentences of particular languages like English are not truth-evaluable, don’t ‘encode’/express any proposition. This is consistent with Claim 1, but it only tells us what ‘linguistic semantics’ ISN’T.

¹ I am grateful to Jay Atlas, Diane Blakemore, Phil Carr, Marjolein Groefsema, Ruth Kempson, Begona Vicente, Deirdre Wilson, David Young—and especially to Robyn Carston—for discussion and/or comments on previous drafts. I gratefully acknowledge an AHRC Research Leave award in support of this work.

On the other hand, when RT offers more positive hints regarding ‘linguistic semantics’, these seem to undermine the distinction, and thus Claim 1, by attributing to expressions of particular languages properties that seem ‘really’ semantic. Robyn Carston offers a representative RT view:

A: The decoding process is performed by an autonomous linguistic system.... Having identified a particular acoustic (or visual) stimulus as linguistic, this system executes a series deterministic grammatical computations... resulting in an output representation which is the semantic representation, or logical form, of the sentence or phrase employed in the utterance. It is a structured string of concepts, with certain logical and causal properties. (2002: 57)

This attributes conceptual and logical—i.e. real semantic—properties to expressions of particular languages. Equally, Sperber & Wilson (1986:72-3) allow that the semantic representation of the English sentence (a) *She carried it* enters into logical relations. For example, it contradicts (b) ‘No one ever carried anything’. But this cannot be unless at least part of what is encoded in (a) is a truth-evaluable proposition (‘Someone carried something at some point in the past’), false when (b) is true. See my (2005) for further discussion.

RT’s ‘linguistic semantics’, in short, is ‘unstable’ in the sense of Cappelin & Lepore (2005). 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 the 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.

This instability is fully reflected in the following passage from Fodor (1998:9), which Robyn Carston (pc) has quoted to me as representing the RT position:

‘English inherits its semantics from the contents of beliefs, desires, intentions, and so forth that it’s used to express.... Or, if you prefer (as I think, on balance, I do), English *has no semantics*. Learning English isn’t learning a theory about what its sentences mean, it’s learning how to associate its sentences with the corresponding thoughts’ (Fodor’s italics).

This hardly clarifies the issue. Either English ‘inherits’ semantics as per Fodor’s first disjunct—in which case what it inherits (and thus has) is the real semantics of LoT—or (as per the second disjunct) it does not. Which is it to be? It is not clear that RT’s position amounts to either of Fodor’s disjuncts. And its distinction between ‘linguistic’ and ‘real’ semantics hardly helps in resolving the matter.

Given this instability, I propose we admit of just the one—real—variety of semantics. Assuming LoT has semantics so understood, this gives us two options:

Option A: Attribute genuine semantic properties to the expressions of particular languages (e.g. English) as well as to LoT (Fodor’s first disjunct)

Option B: Maintain LoT as THE locus of semantics and deny that expressions of particular languages have semantics at all, of any variety (Fodor’s second disjunct).

Option A is the traditional one of course. But we should at least entertain Option B, notwithstanding its *prima facie* improbability. This is what I propose to do. In aid of this, I will appeal to a conceptual project that I and colleagues have been developing: the Representational Hypothesis (RH). The RH offers a perspective from which it can be seen to be not only unnecessary but incorrect to suppose that expressions of particular languages have semantics. Here I can only sketch the idea briefly.² Presented so baldly, it will seem rather startling. From an unlikely angle, it points in the direction of the strongest of Recanati’s (2004:Ch. 9) forms of Contextualism, ‘Meaning Eliminativism’. It follows from the RH that English (for example) indeed has no semantics. So it can’t inherit semantics from LoT. Nevertheless, I believe it follows naturally from the RH that English will—though only to fluent speakers of English—*seem* to inherit semantics from LoT. Seeming-to-inherit-and-to-have and actually inheriting-and-having are different and call for different theoretical models.

Having presented the RH, I will explore the semantic instability in RT by pointing up a crucially related instability in its notion of ‘encoding’.

2. The Representational Hypothesis.

A supposed truism of linguistic theory—Saussurian and Chomskian at least—is that it is necessary to posit entities having properties relating to both sound and ‘meaning’. In Minimalism for example it is assumed that, when a lexical item is selected by the syntactic computation, what’s selected is a (Saussurian) object constituted both by syntactico-semantic and phonological/phonetic properties. At a point in the computation (called

2. See my (2000), Burton-Roberts & Carr (1999), Burton-Roberts & Poole (2006) and (to appear), Chng (1999), for more detail.

‘Spell-Out’), the phonological properties of expressions are stripped out and fed to the phonology. This takes them to PF (Phonetic Form), the interface with articulatory-perceptual systems. What’s left in the computation continues on to LF (Logical Form), the interface with the conceptual-intentional system (LoT).

This—and the general idea of ‘sound with a meaning’—is Chomsky’s ‘DOUBLE INTERFACE’ view of expressions and the computation (1995:2). The computation compositionally projects sound-meaning correspondences encoded in lexical items onto an array of more complex expressions. This double-interface view is thought to be necessary if linguistic communication is to be possible (Chomsky 1995:221). What has syntactico-semantic properties, it is assumed, must also have phonological properties if it is to be utterable —‘tokened’ or ‘realised’—in the acoustic medium of speech (and of course, if the particular language is to be acquirable/learnable on the basis of incoming evidence).

I have come to regard this picture with suspicion. I’ll mention here just two reasons for questioning it. The first is this. If the computation—as a double-interface system—compositionally projects lexical sound-meaning correspondences (Saussurian signs), we should expect it to preserve those correspondences and thus expect isomorphism of phonology and syntactico-semantics in the complex expressions it generates. But, as is well known, this expectation is comprehensively defeated—by ‘the fact that objects appear in the sensory output in positions “displaced” from those in which *they* are interpreted’ (Chomsky 1995:221-2, my emphasis). The very oddity of this way of expressing the matter reflects the problem posed by the double-interface assumption. If, as Cormack & Smith (1997:224) put it, ‘a lexical item does not necessarily appear at Spell Out as a Saussurian sign’, we need to question the existence of lexical items as double-interface objects (Saussurian signs). See Burton-Roberts & Poole (to appear) for detailed discussion.

A second reason for questioning the Saussurian (double-interface) sign is related but is even more fundamental. It concerns arbitrariness. The arbitrariness of the Saussurian sign is universally acknowledged. What is not addressed is WHY it should be arbitrary. It is arbitrary because it is a relation between things that, ultimately, are SORTALLY distinct (Thomason 1972): acoustic phenomena and conceptual structures. This sortal incommensurability is acknowledged in Minimalism, by its assumption that what is PF-interpretable is not LF-interpretable and conversely. Given this sortal rationale of the

arbitrariness, surely, no single entity could possibly be constituted by both sets of properties. But Saussurian signs—lexical items—consist precisely in the conjunction, within a single object, of both sorts of property. For this reason too we should question the very possibility of the Saussurian sign, the double-interface notion of linguistic expression.³ Can we really think of expressions figuring in or generated by the cognitive syntactico-semantic system as utterable, tokened in acoustic phenomena? Acoustic phenomena, I have suggested (2000:44-46), are tokens of acoustic types, not of syntactico-semantic types. Syntactico-semantic types, surely, are tokened in and only in syntactico-semantic structures. On these terms, syntactico-semantic expressions cannot be regarded as ‘realized’ or tokened in speech. We need some other way of modelling ‘sound-meaning’ relations.

The central idea of the RH is very simple: speakers produce phonetic (or other perceptual) phenomena in aid of REPRESENTING expressions manipulated/generated by the syntactico-semantic computation. Ultimately, this boils down to the intuitive and unoriginal idea that speakers utter sounds as a way of perceptually representing—and thus communicating—their thoughts. But, from the RH’s reconstruction of this uncontroversial idea, a picture emerges that contrasts sharply with standard generative assumptions. To bring this out, I need to discuss what I mean by ‘represent(ation)’.

‘Representation’ is not intended here in the sense usual in linguistics, where “‘representation’” is not to be understood relationally, as “‘representation OF’” (Chomsky 2000:159-60). In Chomsky’s ‘non-relational’ sense, what-is-represented is not distinct from the representation itself. A ‘syntactic representation’, for example, is not a representation OF anything; it simply IS the syntax.

In the RH, by contrast, ‘representation’ IS intended ‘relationally’: a representation emphatically is a representation-OF something—something ELSE. To emphasise this, I use the term ‘M-representation’. ‘M’ stands for Magritte, a reminder of his painting *La Trahison des Images*, in which the image of a smoker’s pipe is accompanied by the warning ‘Ceci n’est pas une pipe’. This points up the *distinction* between a representation (or Peircian sign) and what it is a representation (or sign) OF. Simply, a representation-of-a-pipe is NOT (and does not include) a pipe. The RH takes seriously C. S. Peirce’s counter-

³ Jackendoff (1997: esp. 41, 83) is in part motivated by essentially the same thought.

Saussurian stricture that ‘a sign must be OTHER than its object’. (By contrast, the Saussurian sign—as [*signifiant+signifié*])—is partly CONSTITUTED by what it is a sign of, namely *signifié*). On both sortal and semiotic grounds, the RH distinguishes fundamentally between a phonetically constituted M-representation R on the one hand, and, on the other, the structured conceptual (syntactico-semantic) object that R is an M-representation of. Viewed from this perspective, the traditional double-interface view conflates *representans* and *representatum*: by including phonology—along with structured semantics—within the generative system, it conflates facts concerning WHAT-is-represented with logically and empirically independent facts concerning HOW-it-is-represented in the acoustic medium.

I can now explain my earlier scare quotes round ‘sound-meaning relation’. In that phrase, the relational term ‘meaning’ suggests that whatever-it-is-that-sounds-relate-to is to be thought of as a property-of the sound (cf Chomsky’s ‘sound with a meaning’). The RH, by contrast, seeks to emphasise that what the sounds relate to is an entirely independent object, with an independent rationale. It is in fact a conceptual-intentional (C-I) structure, an object not defined by (and innocent of) the fact that sounds relate to it—and emphatically not a property of sound. It is the C-I structure that HAS semantics; it is not the case that the C-I structure IS the semantics-OF anything (let alone sounds).

Let me emphasise here that I am not denying that relevant sounds ‘have meaning’ (for someone). I am denying they HAVE SEMANTICS—a quite different matter. The notoriously vague term ‘meaning’ covers both ‘significance’ and ‘semantics’. All sorts of things can have significance (for someone)—and thus have meaning—without having semantics: black clouds, red litmus paper, blue hands, green lights, ringing bells, raised eyebrows, her failure to return his calls....

Consider again the quote from Fodor. The significance—‘meaning’—of utterances in English consists in the fact that (in the mental life of English speakers) they *function as M-representations* of the kinds of contents that Fodor refers to (thought contents). Given that, speaker-hearers of English will inevitably ‘read into’ (or project onto) a given phonetic m-representation what (they assume) it is an m-representation of. It’s only natural for them to do so. But it does not follow from this that English actually ‘inherits’ the *semantics* of LoT. The closest we can come with any accuracy to saying this is to say that—to English speakers—English will *seem* to inherit the semantics of LoT. But why drag in the technical

notion of ‘semantics’ here? It is enough to say that, for ordinary speakers (of English), the ‘meaning’—i.e. ‘significance’—of utterances in English derives from the thoughts they are used to express. It is the thoughts themselves that have semantics. Speakers get to ‘express’ those thoughts by producing phonetic m-representations of them (even if the m-representations themselves radically undermine the thoughts).

The RH, then, identifies the generative computation as dealing in purely syntactico-semantic properties—recursively hierarchical and non-linear. Conversely, it identifies utterance phenomena as purely phonetic (non-recursive, linear). We don’t utter/hear expressions generated by the computation; what we utter/hear are merely PHONETIC M-REPRESENTATIONS OF those expressions. Rather than (realizationally) EMANATING FROM the generative computation, as in traditional generative thought, PFs are (M-representationally) TARGETED AT that computation. Since the respective properties of *representans* and *representatum* are sortally distinct, the relation is, of necessity, wholly conventional (‘symbolic’ for Peirce, ‘arbitrary’ for Saussure).

As a simple illustration, consider NEGATION. Tradition has it that the English word *not* is a Saussurian double-interface object, constituted by both phonological/phonetic and syntactico-semantic properties. Its semantic properties are those of the logical operator. The RH, by contrast, distinguishes sharply between the operator and the English word *not*. The operator (a) is logically-but-not-phonetically constituted, the word (b) is phonetically-but-not-logically constituted. The two are DISTINGUISHED not only sortally but also in how they are RELATED: conventional M-representation OF (a) BY (b). There is just one such operator but (inevitably, given conventionality) many different phonetic m-representations of it in the speakable languages of the world.

That, distilled, is the representational idea. And here, also distilled, are a couple of its more radical implications, as they appear to me. The first concerns the relation between Chomsky’s Human Faculty of Language (HFL) and LoT. For Chomsky, HFL is recursive, interpretable in conceptual-intentional terms, invariant across the species, innate, and wholly mind-internal. Now these are generally agreed to be the properties of LoT itself. However, as conceived of by Chomsky, HFL must be distinct from LoT because it’s assumed to include phonology (it being, supposedly, a double-interface system). Here’s the implication then. Having excluded phonology FROM HFL—on grounds of its M-

representational relation TO HFL—the RH raises the question why HFL (thought of as a ‘real object of the natural world’—Chomsky 1995:11) and LoT should not be identified. That they are one and the same system is the most parsimonious assumption. And Chomsky’s claims about HFL seem most clearly sustainable if it and LoT are indeed the same. For these reasons (and reasons of space), I’ll assume without further ado that they should be identified, sometimes calling the result ‘L’.⁴

The second implication is really a cluster of implications—for ‘phonology’ and ‘particular languages’. It is, or should be, uncontroversial that a phonological system is a system that determines what counts as a well-formed phonetic string. Now, for the RH, relevant phonetic strings are M-representational. A phonological system, then, is a system that determines what counts as a well-formed phonetic M-representation. Of course, what counts as well-formed M-representational phonetic string differs from language to language. It depends on the particular language—i.e. on its particular M-representational conventions. Now, arguably, a particular language just IS the set of its particular conventions. Accordingly, the RH identifies particular languages AS phonological systems. More specifically, each particular (spoken) language is a distinct, phonologically constituted, Convention System for the Phonetic (M-)Representation of the same single system, L—a $CSPR_{(L)}$. (See Burton-Roberts & Poole (to appear) for further discussion of implications of this idea for ‘phonology’, cross-linguistic variation and what generally passes for language-particular ‘syntax’).

The RH insists that a particular language’s M-representations have only phonetic properties. The (sortally incorrect) theoretical assumption that they also have syntactico-semantic properties arises from theorists agreeing with naïve speakers in projecting onto the M-representations what pertains only to what-they-are-M-representations-of (namely properties of L). Consider ‘parsing’ in this connection. Phonetic M-representations of L require parsing—and parsing is fallible (not ‘deterministic’ as Carston suggests in Quote A above)—precisely because they *don’t* possess syntactico-semantics. Parsing is a matter of

⁴ Bickerton (1996:73-4, 107-112), too, argues for this. However, it is difficult to sustain if you hold—as Bickerton (with Chomsky) does—that HFL is INSTANTIATED in particular languages: for, if HFL = LoT, then LoT must be thus instantiated—surely untenable (Pinker 1992). For the RH, by contrast (see below), the relation of particular languages to HFL/LoT is M-representational. An M-representation-of-*X* is NOT an instantiation-of-*X*.

putting (a) what LACKS syntactico-semantic structure into correspondence with (b) what HAS such structure—on the assumption that (a) was produced with the intention of M-representing (b).⁵

The RH offers a fleshing-out of Fodor’s contention that ‘learning English isn’t learning a theory about what its sentences mean [i.e. their semantics—NBR], it’s learning how to associate its sentences with the corresponding thoughts’. What’s learnt is a system of (phonologically constituted) conventions for the phonetic M-representation of the syntactico-semantic properties of LoT.

3. Varieties of ‘encoding’ in Relevance Theory

What is want to do here is explore Relevance theory’s notion of ‘encoding’ and thereby show that the Representational Hypothesis is considerably more consistent with the quote from Fodor than RT itself is.

The RH was developed in response to problems at the syntax-phonology interface in Minimalism. But it is consistent with—indeed implies—RT’s claim that the Language of Thought (‘L’) is the sole locus of semantic properties. I think the RH clarifies that claim and obviates the ‘real’ vs. ‘linguistic’ ‘semantics’ qualification. It acknowledges the intuition that the phonetically constituted expressions of particular languages have—in some loose sense—‘meaning’ on occasions of use, but without attributing to them objective semantic (conceptuo-logical) properties. Construed as M-representational of what does have such properties, they are indeed SIGNIFICANT (for the construer, at least). If this distinction seems over-subtle, I show below that it suggests a solution to a not-so-subtle problem in ‘lexical semantics’.

Possibly, the tension surrounding ‘linguistic semantics’ arises from RT’s reluctance to abandon the Saussurian/Chomskian ‘double interface’ tradition despite RT’s view of LoT as the sole locus of semantics. This tension, I believe, is reflected in RT’s ‘encoding’.

Carston (2002: 57-58) refers to David Lewis’ (1970/83) distinction between the ‘genuinely semantic’ (conceptual-intentional) relation between some symbolic system and the world of NON-symbols on the one hand and, on the other, the relation of ‘encoding’.

⁵ Dynamic Syntax (Kempson et al 2001) doesn’t appeal to (M-)representation, but its process of building structure on the basis of string information is consistent with the representational idea. With the RH, Dynamic Syntax effectively rejects the double interface assumption.

Encoding deals merely with (non-conceptual-intentional) relations between one symbolic system and another. RT's 'linguistic semantics' refers to a system of utterable symbols and its encoding relation to L. As with Morse code, all that is needed or appropriate here is not a SEMANTICS for the encoding system but a set of CODING CONVENTIONS. On these terms, particular languages don't have semantics, they are mere codes.⁶

The force of Lewis' (important) distinction rests crucially upon a further distinction, however: that between the ENCODING with WHAT IS ENCODED. To maintain Lewis' encoding-semantics distinction, we must be careful not to conflate the encoding and what is encoded. To do so would risk attributing the genuine ('real') semantic properties of LoT to mere encodings of them in the (utterable) symbols of particular languages. If we strictly maintain these distinctions, then no remotely semantic (conceptual-logical) properties can be attributed to the utterable encodings of a particular language.

All this suggests something similar, even identical, to the Representational Hypothesis. 'Encoding' on these terms is tantamount to the RH's notion of M-representation. On that assumption, I'll call it 'M-ENCODING'.

All this notwithstanding, RT generally operates with a quite different notion of 'encoding', inherited from the double-interface tradition. That is, when it characterizes an expression *E* as 'encoding' some conceptual-logical property *f*, RT attributes *f* as a property-of *E*. I will call this notion of encoding 'C-ENCODING'—'C-' for 'constitutive'. C-encodings are CONSTITUTED by what they encode (as well as by phonetic properties—i.e. by how they encode it). This C- notion of encoding completely undoes the appeal to Lewis, undermines Claim 1 and gives rise to a range of problems which I will illustrate by reference to negation, narrowing, and numerals.

3.1. Problems with C-encoded negation.

Carston (2002: ch. 4, esp 311) conceives of *not* in C-encoding terms: for her, the word HAS the logical properties of (and thus is) the operator that it 'encodes'. The first problem with this is that, if *not* C-encodes (rather than M-encodes) the logical operator, then we have to accept that the (real) operator itself—a function from true(false) to false(true)—figures in

⁶ Thanks to Deirdre Wilson for discussion.

the domain of ‘linguistic encoding’. This is anomalous given RT’s claim that the ‘linguistic encoding’ is a NON-TRUTH-THEORETIC domain.

Furthermore, if the English word *not* C-encodes—and thus is—a logical operator, there must be scope-of-negation WITHIN THE LINGUISTIC ENCODING. Where real negation is, there also is (must be) scope of negation. Against this, it is often unclear from the linguistic encoding what the scope of negation is. This is suspicious: were the encoding such as to C-encode and thus include the logical operator itself, and thus have a genuine logical form, we should expect it to wear scope-of-negation on its sleeve. We can’t say the encoding is scopally ambiguous if it C-encodes (and thus has) a logical form. Logical form, by definition, is *not* ambiguous. To admit that the linguistic encoding could be semantically ambiguous would be to concede that the linguistic encoding and logical form are distinct—that the latter isn’t a property of the former—in short, that logical form is NOT C-encoded. Ambiguity indeed is a central plank in RT’s argument that truth-theoretic properties are ELSEWHERE (in the C-I structures of LoT).

But for Carston (2002: Ch 4), negation—and hence scope-of-negation—is C-encoded. And, notwithstanding the in clarity of the encoding, she wants to avoid scopal ambiguity. How to square this circle? She does it by insisting that (C-)encoded negation always has wide scope. When this doesn’t square with the thought communicated, she argues, narrow scope is derived by pragmatic inference. Well, this does the trick, but it is somewhat stipulative. The point is that cases in which this strategy is available are precisely cases in which there IS nothing in the linguistic encoding that indicates scope. And RT is usually the first to insist that, in such cases, there is nothing ‘encoded’ (i.e. C-encoded).

That scope of negation is ALWAYS pragmatically inferred is Jay Atlas’ (1989, 2005) scope neutrality thesis. For Atlas, there IS no scope of negation in the linguistic encoding (the sentence). Carston writes that the ‘radical indeterminacy’ implied by Atlas’ proposal is ‘completely at one with the basic thesis of this [her] book’ (2002:287). Nevertheless, she doesn’t adopt that proposal—and it is indeed difficult to see how it could be adopted if the word *not* is a linguistic C-encoding of, and thus HAS, logical properties. To repeat, where negation is, there also is scope-of-negation. Although Atlas (1989) opens with a suggestive and relevant discussion of representation (effectively—for him as for me—‘M-representation’), he doesn’t invoke it in connection with negation. He too assumes C-

encoding rather than M-encoding here. At least, he never suggests that *not* doesn't have logical properties (see also Atlas 2005), indeed refers to it (1989:80) as 'one of our language's most basic logical words'. And, while assuming that semantics IS semantics-of-sentences (sentences of English for example), he gives no account of what his scope-neutral sentence-semantics actually consists in. See my (1991:3.5) and Carston (2002:287-8) for discussion.

I believe that M-encoding—M-representation—offers a framework within which Atlas' proposal and Carston's (and RT's) 'basic thesis' can be properly realised. Operating with M-encoding, we can—must—say that the encoding includes nothing with logical properties, but merely phonetic material (*not*) which, by the conventions of English, is used to M-represent (M-encode) something else, namely a logical operator located in and only in LoT. 'Scope neutrality' follows automatically—or something stronger since, if the operator simply isn't there in the (M-)encoding, it's beside the point to say the encoding is 'sense general' or 'neutral' with respect to scope. One might as well say that cabbages or kings are 'scope-neutral'.

On these terms—and as Atlas suggested—pragmatics is required, not on SOME occasions nor in order to narrow C-encoded wide-scope negation, but on ALL occasions, to infer the logical properties of the thought intended to be communicated by some—ontologically purely phonetic, but functionally M-representational—utterance.

It might be thought the RH is committed to the linguistic M-encoding being scope ambiguous. In fact this makes no sense in RH terms, if 'ambiguity' is understood (as it standardly is) as a SEMANTIC property. To repeat, it follows from the RH that linguistic encodings—as M-encodings, or M-representations—DON'T HAVE semantic properties. So they can't be semantically ambiguous. This is not to say they can't be subjectively 'ambiguous'. But that simply means that it—scope of negation in thought, or whatever—can be UNCLEAR to a particular hearer on a particular occasion. Subjective 'disambiguation' is a matter of the hearer escaping this subjective state by ascertaining—to his own satisfaction at least—what the speaker intended. This is not what linguists generally mean by 'semantic ambiguity'—which, as objective, is supposed to remain in

the absence of subjective ambiguity.⁷ The very possibility of subjective ambiguity arises precisely because linguistic encodings, as purely phonetic, are—as a matter of ontological necessity—M-representationally indeterminate with respect to the thoughts they are intended to evoke.

3.2 *Loosening and narrowing*

I am suggesting that, while RT operates (in the light of Lewis' semantics-encoding distinction) with M-encoding at a general level, when it comes to the analysis of particular phenomena (e.g. negation) it in fact operates with C-encoding.

It is clearly C-encoding that underlies the distinction between pragmatic narrowing (enrichment) and loosening (broadening)—Carston (2002: ch. 5). That is, it is assumed that there are two conceptual domains: (a) C-encoded in words—'lexical concepts'—and (b) in LoT, actually entertained in thought. Loosening and narrowing are involved in getting from conceptual domain (a) to conceptual domain (b). Take loosening first. Examples are *bald* and *raw*. These words are assumed to C-encode the non-gradient, absolute concepts HAIRLESS and UNCOOKED respectively. However, it is not generally the lexical concept that's actually entertained in thought but one of a range of distinct, looser, gradiently related concepts.

Now it might be argued that it is empirically incorrect (i) to pick the non-gradient concept as the (C-)encoded concept and posit loosening when necessary, rather than (ii) pick a gradient concept and positing narrowing when necessary. My contention, rather, is that it is arbitrary. As Robyn Carston (pc) notes, the availability of *He wasn't strictly speaking bald* suggests (i). Against this, the non-tautological feel of *completely/totally bald* suggests (ii). Does the acceptability of the phrase *very bald* indicate that the lexical concept is gradient or that it is non-gradient but can be used loosely? I don't think this is empirically decidable. But, given C-encoding (and thus a domain of 'lexical concepts'), there's no escape: we must decide.

⁷ If *bank* is ambiguous as a matter of C-encoded (semantic) fact, it should be ambiguous regardless of context. But everyone concedes that 'ambiguity' is generally resolvable in context—in other words, that ambiguity is context-dependent. This suggests that, when semanticists/pragmaticists purport to be addressing an objective, semantic phenomenon, they are in fact addressing a subjective, pragmatic phenomenon. See my (1994).

More on loosening below. It is really narrowing that concerns me here. Carston assumes that the English word *tired*, for example, C-encodes a single, very general, highly abstract concept. In thought, by contrast, there is an ‘virtually indefinite range of finely distinct [highly specific] concepts’—from slight lassitude to total exhaustion, from [ENOUGH TO MAKE WALKING THE DOG MORE ATTRACTIVE THAN JOGGING] to [IMPOSSIBLE TO MAKE IT TO BASE-CAMP THOUGH THAT’S THE ONLY CHANCE OF SURVIVAL]. This disparity—between the single, highly general, lexically C-encoded concept and the multiplicity of highly specific concepts in thought—calls for narrowing.

Carston vividly delineates a fundamental problem in connection with narrowing: in many cases, the lexical semanticist will need to posit ‘lexical concepts’ that are SO extremely general and SO abstract, it is difficult to imagine that any actual thought COULD involve them. *Open* is a classic case. As Carston notes, we use this verb to talk of John ‘opening’ a window, his mouth, a book, a briefcase, the curtains, a package, a wall, a wound. Think also of: his fingers, his dog’s mouth, his bowels, a drawer, a bottle, the throttle, a computer file, the bidding, a chess game, a shop (for the first time, or as usual at 9.00am) or Parliament. As Carston puts it:

B: “[W]hen we try to think about the general concept OPEN and to have a thought in which such a general concept features, as opposed to any of the more specific concepts that we grasp,...the experience is an odd one, as we seem to have no definite thought at all.... [W]hat is not at all clear is whether we actually have (hence sometimes try to communicate) thoughts in which this very general lexicalised concept features as a constituent....But surely if the word *open* encodes a concept we should be able to have thoughts which include that concept...” (2002:361-2)

Exactly. When an English speaker says that Max opened a book, it’s not part of what she says or means that Max did anything in common with the Queen opening Parliament, Kasparov deploying the Queen’s Pawn Gambit, my use a corkscrew or.... The putative C-encoded lexical concept doesn’t just envelop us, as analysts, in a cloud of unknowing—it’s hard to accept that it figures even in any ordinary use of the word.

An admittedly strong intuition underlying the search for a unitary semantics for *open* might go like this: ‘Surely, since we use the same word for all those activities, they must have SOMETHING in common—and, surely, that “something” constitutes the semantics of

the word.’ We need to resist this. Everything has something in common with everything else but we don’t use the same word for everything. More significantly, not all languages use one and the same word for exactly the range of activities that in English we use *open* for (Bowerman & Choi 2003). That intuition is in danger of leading to a ‘rightly-are-they-called-pigs’ school of lexical semantics. Like the Argument from Design, it gets it round the wrong way. What all and only those activities have in common is that—in English—we use the same word for them.⁸

3.3. *The problem of N.*

The problem—the fugitive character of lexical semantics—arises with numerical expressions. In use, *N* can mean any of

- (a) [EXACTLY *N*]
 - i. Q: How many children do you have? A: **Three**.
 - ii. I had five and gave three to Ed. So I have **two** left.
 - iii. A triangle (but not a quadrilateral) has **three** sides.
- (b) [AT LEAST *N*]
 - i. You must be **eighteen** to vote.
 - ii. If you have **four** children, you qualify for benefit.
- (c) [AT MOST *N*]
 - i. I must pare that article down to **sixty** pages.
 - ii. We’re allowed **thirty** days’ holiday a year.

The favoured (neo-Gricean) analysis is that *N* is lower-bounded semantically—has (b) as its (C-encoded) semantics—and upper-bounded by implicature (Horn 1989: Ch. 4). This pragmatically derives (a) from (b). However, as Carston (1988) notes, this doesn’t

⁸ Consider Denham’s ‘Lines on the Thames’ (quoted by Richards 1936:121):

*O could I flow like thee, and make thy stream
 My great exemplar, as it is my theme!
 Though deep, yet clear; though gentle, yet not dull;
 Strong without rage; without o'erflowing full.*

As Richards notes, Dr Johnson praised this on the grounds that ‘the particulars of resemblance are so perspicuously collected’—displaying what for Aristotle was ‘an eye for resemblances’. But this assumes the lines are an exercise in natural history, reporting on supposed resemblances (between a river and a mind) out there, rather than a work of literary—that is, linguistic—art. The resemblances hold only ‘under a description’.

accommodate (c). Conversely, were *N* semantically upper-bounded, as (c), and lower-bounded pragmatically, yielding (a), this wouldn't accommodate (b). Sadock (1984) favoured (a) as the semantics, suggesting that (b) and (c) arose from loose use of *N*, by analogy with 'France is hexagonal'. But, as Carston notes, loose use of *N* is quite different from (b)-(c). The true analogy is between 'France is hexagonal' and examples in which *N* is used to convey [APPROXIMATELY *N*], which corresponds to none of (a)-(c).

On the basis of such considerations, Carston (1998) suggests that the semantics of *N* (what *N* C-encodes) cannot be any of (a)-(c). For her—and Atlas (1992)—its semantics is more general than any of those. But, as with *open*, the question is: what could that semantics possibly amount to? Logically, the inclusive disjunction $[[\text{AT LEAST } N] \cup [\text{AT MOST } N]]$ might do the trick. However, possibly because it comes too close to polysemy, Carston doesn't consider it. And Atlas (1989) explicitly rejects a disjunctive analysis of sense generality.

Instead, Carston's account of the concept C-encoded by *three*, for example, is '[X [THREE]]', where the value for 'X'—chosen from {[AT LEAST], [AT MOST], [EXACTLY]}—must be supplied pragmatically (a case of saturation?). Now, arguably, this just transfers variability in the (putative) semantics of *three* onto 'X'. And there is a more serious problem. On C-encoding terms, we're looking for a SEMANTIC DEFINITION of *three*—and, considered as such, '[X [THREE]]' surely won't do. Independently of, or prior to, assigning a value to 'X', we still have '[THREE]'. Notwithstanding the typographical distinction (small caps), '[THREE]' can hardly figure (non-circularly) in a semantic definition of *three*.

C-encoding commits us to identifying a SINGLE ('lexical') concept consistent with a VARIETY of uses. But, again, even when—as with *N*—the variety isn't huge, we seem to be facing a concept that's quite inscrutable. This is the general problem posed by the (C-encoding) assumption that a particular language constitutes a conceptual domain (of 'lexical concepts') distinct from and in addition to the conceptual domain constituted by LoT.

I return to *three* below, where I pursue Jerry Sadock's more recent observation that 'If sometimes the cardinal numbers are upper bounded by implicature and sometimes lower bounded, we are lead to conclude that they have no semantic content at all'(2005)—and the idea that (therefore) implicature can't be involved in how uses of *three* are understood.

4. Carston's response.

In responding to the problem (in connection with *open*), Carston comes tantalisingly close to M-encoding, the M-representational relation of the RH.

C. Could it be that the word...does not encode a concept, but rather 'points' to a conceptual region...in memory? (360).

As I interpret this, Carston is here (effectively) questioning the notion of C-encoding and thus entertaining the possibility of M-encoding. Indeed, the pointing idea (borrowed from Sperber & Wilson 1997/8:196-7) exactly captures the Magrittian/Peircian representational distinction I'm after. A pointer, [α], is OTHER THAN what-it-points-to, [β]. [α] does not, in virtue of pointing to [β], partake of any of [β]'s properties (conceptual, in our case).

In the event, however, it appears that M-encoding—and thus the denial of any conceptual property (semantics) located in utterable words—is not what Carston has in mind. Instead, she replaces talk of 'encoding a concept' with talk of 'encoding a concept schema or pro-concept' (363). I deal here just with 'concept schema', returning to 'pro-concept' below. She writes:

D. Suppose it is right that there is a sizeable class of words that do not encode particular concepts (senses) but rather concept schemas or pointers, or addresses in memory (which of these is the best metaphor remains unclear)...(363)

It seems that, for Carston, 'pointer' is just another metaphor along with 'concept schema'. But these—for me, at least—are utterly different. To repeat, using a phonetic pointer to (phonetically M-encoding/M-representing) a conceptual region in thought doesn't require us to attribute ANY conceptual property to the word (the pointer) itself—quite the reverse. Carston's notion of 'concept schema', by contrast, seems called for precisely by the perceived need to attribute SOME kind of conceptual property to the word—a need arising only on C-encoding assumptions. Carston herself (363) points to the parallelism between 'concept schema' and her 'schema or template for a range of propositions'. But the latter, I've suggested (2005:396), precisely IS a proposition, albeit a

very general one. By the same token, I suggest a ‘concept schema’ just IS a very general, highly abstract concept.⁹ Despite the change in terminology, this brings us right back to the problem we started with (quote B above)—and, furthermore, the task of specifying that schema’s content, and thus offering a substantive semantics for *open* (one that goes beyond Carston’s ‘[OPEN]’).

Whatever we call them, on these C-encoding terms, we are going to have to admit of what Carston (very honestly) describes as ‘a whole additional population of mental entities’ distinct from the concepts actually entertained or entertainable in thought, which ‘don’t seem to have any function in mental life except to mediate the word/concept relation’ (363). Furthermore, we face

- E.** [a] challenging question that arises concerning acquisition: if word meanings are these abstract schematic entities that do not feature in our thinking about the world, how do we ever manage to acquire (learn) them?” (363)

She writes

- F.** There must be some process of abstraction... 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 (364).

Carston is persuaded (in contrast to Hintzman’s (1986) multiple trace memory model) that the multiple traces left by experience of particular uses resolve themselves into a distinct general conceptual representation (rather like those multiple images that, seen from a distance, compose a picture of the Mona Lisa). The difficulty is not so much how the abstraction arises—assuming that it does—but what cognitive function it serves. Its function couldn’t be to enable ‘understanding of subsequent utterances’, precisely because

⁹ The alternative interpretation—that, in positing ‘concept schemas’, Carston is indeed denying that words have conceptual properties—would leave the notion of ‘concept schema’ quite unclear. And, in fact, on p. 375, responding to Recanati (1998), she makes it clear she is not pursuing ‘the radical claim that there is no lexical meaning in the sense of stable encoding’ but a ‘more conservative one... on which words do encode something, albeit something very schematic...’.

it is *post hoc*—it arises FROM, and presupposes, a prior understanding of utterances. (Hintzman 1986, esp. 422-3).

5. Biting the M-representational bullet.

Carston has articulated a genuine problem, implied by the very idea of ‘lexical’ or ‘linguistic semantics’. Given these problems, and our overarching endeavour of identifying LoT as the sole locus of conceptuo-logical (semantic) properties, I suggest we unequivocally jettison C-encoding in favour of M-encoding, deny that words have conceptual properties, and thereby deny the existence of ‘lexical semantics’—on the grounds of the Magrittian/Peircian DISTINCTION between properties of the phonetic *representans* (the word) and those of the conceptual *representatum*.

What we have in *bald*, *tired*, *open*, *three* is, in each case, a single *representans* and, quite distinctly, a multiplicity of conceptually related *representata* in thought, on occasions of their use. The singularity we have in each of these is not conceptual but phonetic. The multiple traces resolve themselves round a PHONETIC attractor (or ‘address’). This disposes of the problem of identifying a single, context-invariant, cognitively suspect (‘lexical’) concept as the ‘semantics of the word’.

A consequence is a Fodorian emptying of the lexicon (Fodor & Lepore 1998) and, I believe, a dismantling of the distinction between logical entries for words and encyclopaedic knowledge/presuppositions (Searle 1980). This resolves a tension in Carston’s work in this connection: while explicitly endorsing (2002: 141, 321) Fodor’s arguments against lexical decomposition (essentially, arguments against the analytic-synthetic distinction), Carston (with Sperber and Wilson 1986:86) nevertheless posits, in addition to encyclopaedic entries for words, logical entries that ‘capture certain analytic implications of the [lexical] concept’. See Groefsema (this volume), Vincente (2005) for further discussion.

In fact, the M-representational approach goes further than Fodor. Fodor has it that WORDS have denotations (surely not consistent with the Fodor quote above). I suggest it is, at best, particular USES of words that have denotations (Recanati 1998). And even then, what I really mean is that it is the CONCEPTS (in thought—where else?) M-represented by such word-uses that in fact have denotations.

5.1. 'Three' again.

On these terms, rather than look for a semantic definition of the word *three*, we might attempt to say something about the range of concepts that *three* can be used to point to, or M-represent.

On some occasions, the concept actually entertained in thought and M-represented by the use of *three* is [THREE OR MORE]. And the concept [THREE] that figures there is [EXACTLY THREE]. This gives [EXACTLY THREE OR MORE THAN EXACTLY THREE]. Those are occasions on which *at least three* would be a more precise M-representation. On other occasions—when *at most three* would be more precise—the entertained concept is [EXACTLY THREE OR LESS THAN EXACTLY THREE]. And on others again, it is [EXACTLY THREE].

Given the distinction between the M-representation and the concept it M-represents—so we're not talking of any conceptual property of *three* itself—the circularity of using '[THREE]' to explicate a putative semantics for *three* doesn't arise. It is the concept [(EXACTLY) THREE] that has semantics. Its semantics is constituted by its denoting a number—the number of dots in a 'therefore' sign (\therefore), for example.

Notice that this captures the unity of uses of *three* (i.e. what all uses of *three* have in common): what those different concepts all have in common is the fact that [EXACTLY THREE] figures in their composition. It is [EXACTLY THREE], then, that defines the conceptual range of (and constrains) the use of *three*. So, presumably, the convention governing the use of *three* effectively says: use *three* to m-represent any cardinal concept in thought in which the concept [EXACTLY THREE] figures.

This proposal has an important implication, which I think could only emerge on an M-encoding/representational approach: it allows us to acknowledge [EXACTLY THREE] as defining the conceptual range of uses of the word—and as the concept referred to in the relevant convention—WHILE SIMULTANEOUSLY DENYING THAT [EXACTLY THREE] IS THE SEMANTIC DEFINITION OF THE WORD. This strikes me as a highly satisfactory result because—even were we to allow that the word has a semantic definition, which I don't—we already know that [EXACTLY THREE] is not (and cannot be) that definition.

Furthermore, the result is intuitive. I suspect that Sadock's (1984) proposal—that *three* actually means [EXACTLY THREE]—would be the intuitive/pretheoretical definition-of-choice for most ordinary speakers. If I am right in that, then the M-encoding approach has the merit of reconciling that pretheoretical intuition with the (later) conclusion of Sadock (2005) that the word in fact has no definition. Indeed, in identifying [EXACTLY THREE] as involved in all uses of *three* while (correctly) denying that it constitutes the definition of *three*, this approach could be said to EXPLAIN that (real but mistaken) intuition.

What all this shows more generally is that identifying the conceptual range of uses of a word *W* is NOT THE SAME AS identifying *W*'s putative semantics. It is possible, then, that the more daunting task of identifying the conceptual range of uses of the word *open* (which I won't attempt) wouldn't anyway yield anything that could be regarded as its semantics.

5.2. *Narrowing-loosening and conceptual-procedural*

Carston (2002: Ch 5) seeks to unify loosening and narrowing as a single process of 'concept adjustment' resulting, across the board, in explicature. This, she suggests, departs from the standard RT (asymmetric) view that only narrowings—as enrichments or 'developments'—are explicated, and that loosening, by contrast, are implicated.

However, if (as Carston assumes) there are two conceptual domains, (a) C-encoded in words and (b) entertained in thought, then—like it or not—some inferences from (a) to (b) will result in 'narrowings' and others in 'loosening'. With C-encoding (and thus 'concept adjustment'), you are not going to disband the loosening-narrowing distinction. In which case, it is difficult to see how the results of ALL such inferences could be explicated. Of course it depends on the criterion for 'explicature' (as against 'implicature'). This is unclear in RT since 'explicature' is defined in terms of 'development' (of the logical form of the sentence uttered) but 'development' itself is not defined—see my (2005). If we go with Carston (1988), we might take a proposition *P* to be a 'development' of the LF—and thus explicated—if *P* unilaterally entails the LF. True, this doesn't cover all cases that RT has claimed to be explicatures (and is not consistent with 'linguistic semantics' being non-truth-theoretic—see my 2005: 398) but it is at least clear conceptually and seems pretheoretically consistent with explicatures being 'enrichments' of the LF.

On these terms, explicatures are conceptually/logically MORE SPECIFIC than what is encoded. But then, only narrowings could count as developments, and thus as explicated, not loosening—just as in the asymmetric account that Carston seeks to reject. What makes ‘development’ and ‘enrichment’ pretheoretically appropriate in describing concept adjustments that result in narrowings—and thus allows us to see narrowings as explicated—is that the narrowed concept INCLUDES AND BUILDS ON the content of the lexical (C-encoded) concept. In ‘loosening’, by contrast, the lexical concept is not included—it is DROPPED AND REPLACED. How then could any concept adjustment that resulted in a loosened concept count as an ‘enrichment’ or ‘development’ of the lexically encoded concept? As long as we HAVE a distinction between loosening and narrowing, we are not going to be able to say, with any plausibility, that concept adjustment results (across the board) in explicature, as Carston proposes.

But, again, it depends on what ‘development’ is—and that’s the problem. The problem is intensified rather than resolved when Carston writes:

G: The characterization of an ‘explicature’ as a communicated assumption which is a *development* of a logical form of the utterance... can be maintained provided that the notion of a ‘development’ of a logical form is understood to include pragmatic adjustments to linguistically encoded concepts which may involve dropping logical or definitional elements of the encoded concepts.
(2002:342, original italics)

I believe the RH can make a positive contribution here. Narrowing and loosening are forms/results of ‘concept adjustment’. The clearest way to disband the narrowing-loosening distinction is to deny the existence of ‘concept adjustment’. This denial follows directly from M-encoding/M-representation. If words aren’t conceptual C-encoders in the first place—but merely phonetic M-representational pointers to conceptual regions in thought—then there IS no lexical conceptual domain (a). So there can’t be any conceptual disparity between (a) and (b). There are, in short, no ‘lexical concepts’ to be ‘adjusted’—or, therefore, either narrowed or loosened.

Instead, and across the board, we have a single inferential process from a speaker’s ‘phonetic effort’ (Searle 1965) to a thought—assuming those efforts were M-representational. Now a speaker’s M-representational ‘phonetic effort’ is what RT calls an

‘ostensive acoustic stimulus’. The inference is from just such a stimulus to a thought—and it makes no sense to say that some but not all of these inferences yield ‘enrichments’ or ‘developments’ of what was ‘encoded’. In fact, it makes no sense—in M-encoding terms—to say ANY of them do. So, if Relevance Theory wants to say that the result of the inference is EVER explicated, we have no reason not to say that it ALWAYS is. Whatever you ‘meant’—whatever you had in mind (!) and were M-representing—in using *tired*, *bald*, *three* or *open*, that’s what you ‘said’, ‘explicated’, ‘expressed’ and were committed to.

In this connection, we might dismember the notion of ‘pro-concept’, discarding the ‘-concept’ bit and developing the ‘pro-’ bit. As is well known, pro-forms (such as *she* and *it* in *She carried it*) are now treated in RT as having ‘procedural meaning’, not ‘conceptual meaning’; they encode procedures that a hearer should engage in to infer explicated conceptual-intentional properties. Recall that it was precisely the attribution of ‘conceptual meaning’ (as against procedural meaning) to the linguistic encoding that called for the qualification embodied in Claim 2 and led to the undermining of Claim 1.

Speculating: if we want to allow that linguistic encoding ever C-ENCODES anything, we might allow that what it C-encodes is procedural. That is, since the phonetically constituted expressions of particular languages function (for those aware of the M-representational conventions of the given language) as M-representations of the conceptual-intentional structures in terms of which thought is couched, it is not unreasonable to think of those expressions as C-encoding (i.e. constituting) instructions to construct such conceptual-intentional structures. In which case, the Representational Hypothesis could be expressed in terms of the conceptual~procedural distinction by saying that the words of particular languages function as CONCEPTUAL M-ENCODERS in virtue of being PROCEDURAL C-ENCODERS. On these terms, all linguistic ‘meaning’ (encoding) is procedural—M-representationally pointing the hearer towards structures in LoT, the unique locus of conceptual-intentional (i.e. semantic) properties.

6. Envoi.

Carston’s discussion raised the general question of how we get to understand other people’s uses of words. For her it is the question of how abstract lexical semantics arises (in a given I-language) and is an ‘acquisition’ question. The position advanced here is that

there are no abstract schematic lexical concepts to ‘acquire’. Whatever is going on here, it is not a process of acquisition and it doesn’t involve abstraction. It seems, rather, to be a process of LEARNING—learning to USE bits of phonetic material in roughly the same ways as others in the community. This, we have seen, presupposes some understanding of the utterances of the community. The circularity that threatens here can be avoided by positing a lively capacity for pragmatic inference and mind–reading (Carston 2002:30).

The ‘representational conventions’ that I’ve been alluding to may be nothing other than—in one sense of ‘conventional’—the conventional (i.e. habitual) use of words in a community. In this sense of ‘convention’, the conventions DETERMINE use precisely because they are themselves DETERMINED BY use (Pelczar 2000:503).

When I started teaching pragmatics, I used to contrast Gricean ideas with the Wittgensteinian slogan ‘Don’t ask for the meaning, ask for the use’, suggesting that Wittgenstein operated in a benighted world with no semantics-pragmatics distinction. Well, perhaps Wittgenstein should rather be thought of as the father of radically radical pragmatics. Certainly Hintzman’s multiple trace theory—appositely cited by Renanati 1998—points towards a Wittgensteinian concept of lexical understanding.

REFERENCES.

- Atlas, J. 1989. *Philosophy without ambiguity*. Clarendon Press.
—1992. ms. ‘Why three does not mean 3’. Pomona College.
—2005. *Logic, meaning and conversation*. Oxford University Press.
Bickerton, D. 1996. *Language and human behaviour*. UCL Press.
Bowerman, M., & S. Choi. 2003. ‘Space under construction: language specific spatial categorization in first language acquisition.’ In D. Gentner and S. Goldin-Meadow (eds.) *Language in Mind*. MIT Press. 387-427.
Burton-Roberts, N. 1991. Review of Atlas 1989. *Mind & Language*. 6:161-176
—1994. ‘Sentence, utterance and ambiguity: a representational approach.’ *Transactions of the Philological Society*. 92: 179-212.
—2000. ‘Where and what is phonology? A representational perspective.’ In N. Burton-Roberts, P. Carr & G. Docherty (eds.) *Phonological knowledge*. OUP. 39-66.
—2005. ‘Robyn Carston on semantics, pragmatics and encoding.’ (Review of Carston 2002) *Journal of Linguistics* 42: 389-407.
— & G. Poole. 2006. ‘Syntax vs phonology: A representational approach to stylistic fronting and verb-second in Icelandic’. *Lingua* 116:562-600.
— & G. Poole. To appear. “‘Virtual conceptual necessity”, feature-dissociation and the Saussurian legacy in generative grammar’. *Journal of Linguistics* 42:3.
Cappelen, H., and E. Lepore. 2004. ‘Radical and moderate pragmatics: does meaning

- determine truth conditions?' In Z. Szabo (ed.) *Semantics vs. Pragmatics*. Oxford University Press. 45-71.
- Carston, R. (1988). 'Implicature, explicature and truth-theoretic semantics.' In Kempson, R. (ed.) *Mental Representations*. 155-81. Reprinted in Davis, S. (ed.) (1991) *Pragmatics: a reader*. Oxford: University Press. 33-51.
- 1998. 'Informativeness, relevance and scalar implicature.' In R. Carston & S. Uchida (eds) *Relevance theory: implications and applications*. Benjamins. 179-236.
- 2002. *Thoughts and utterances: the pragmatics of explicit communication*. Blackwell.
- Chomsky, N. 1995. *The minimalist program*. MIT Press.
- 2000. 'Language from an internalist perspective.' In N. Smith (ed.) *New horizons in the study of language and mind*. Cambridge University Press. 134-163.
- Chng, S. 1999. *Language, thought and literal meaning*. PhD thesis, University of Newcastle.
- Cormack, A., and N. Smith. 1997. 'Checking features and split signs.' *Working Papers in Linguistics* 8. University College London, 223-322.
- Fodor, J. 1998. *Concepts: where cognitive science went wrong*. Oxford University Press.
- and E. Lepore. 1998. 'The emptiness of the lexicon: reflections on James Pustejovsky's *The Generative Lexicon*.' *Linguistic Inquiry* 29. 269-88.
- Groefsema, M. 1997. 'Concepts and word meaning.' This volume.
- Hintzman, D. '“Schema abstraction” in a multiple-trace memory model.' *Psychological Review* 93: 411-428.
- Horn, L. 1989. *A natural history of negation*. Chicago University Press.
- Jackendoff, R. 1997. *The architecture of the language faculty*. MIT Press.
- Kempson, R., W. Meyer-Viol and D. Gabbay. 2001 *Dynamic syntax*. Blackwell.
- Lewis, D. (1970). 'General semantics.' *Synthese* 22. 18-67. Reprinted in Lewis, D. (1983). *Philosophical papers*. Vol. I. Oxford University Press. 189-229.
- Pelczar, M. 2000. 'Wittgensteinian semantics.' *Noûs*. 34:4. 483-516.
- Pinker, S. 1992. Review of Bickerton, D. *Language and species* (1990 Chicago UP). *Language* 68: 375-82.
- Recanati, F. 1998. 'Pragmatics.' *Routledge encyclopaedia of philosophy*. Vol. 7. 620-33.
- 2004. *Literal meaning*. Cambridge University Press.
- Richards, I. A. 1936. *The philosophy of rhetoric*. Oxford University Press.
- Sadock, J. 1984. 'Wither radical pragmatics?' In D. Schrifin (ed) *Meaning, form and use in context*. Georgetown University Press.
- 2005. Abstract for 'Math according to GARP'.
<http://www.lcs.pomona.edu/JayFest/abstract.htm>
- Searle, J. 1965. 'What is a Speech Act?' In M. Black (ed) *Philosophy in America*. Allen & Unwin.
- 1980. 'The background of meaning' in J. Searle et al (eds.) *Speech act theory and pragmatics*. Reidel. 221-32.
- Sperber, D. and D. Wilson. 1986. *Relevance: communication and cognition*. Blackwell.
- 1998. 'The mapping between the mental and the public lexicon.' In P. Carruthers and J. Boucher (eds.) *Language and thought*. 184-200.
- Thomason, R. 1972. 'A semantic theory of sortal incorrectness.' *Journal of Philosophical Logic* 1. 205-298.

Vincente, B. 2005. 'Meaning in relevance theory and the semantics/pragmatics distinction'. In Coulson, S. and B. Lewandowska-Tomaszczyk (eds.) *The literal and non-literal in language and thought*. Petr Lang.