Available online at www.sciencedirect.com
SCIENCE DIRECT®
Lingua 114 (2004) 621–628
Lingua
www.elsevier.com/locate/lingua

Book review

Thoughts and Utterances: The Pragmatics of Explicit Communication
Robyn Carston; Blackwell, Oxford, 2002

To those who have not followed recent advances in pragmatics, the sub-title of Robyn Carston's book may seem surprising, even paradoxical. Indeed, until recently, the dominant view among most linguists and philosophers was that pragmatics dealt with implicit aspects of communication, mainly implicatures, while explicit, literal meaning was delivered by decoding the linguistic (semantic) content of utterances. Grice clearly held that view: even though he recognized that pragmatic processes of disambiguation or reference assignment have to contribute to 'what is said', he saw this sort of contribution as very limited and peripheral. In this book, Carston explores an alternative view, according to which words merely evoke (rather than directly encode) thoughts—hence even the computation of explicit, literal meaning relies extensively on pragmatic-inferential processes. The result is a fascinating study of how semantics and pragmatics conspire to enable humans to convey long and complex thoughts through often short and simple linguistic utterances.

Carston starts out (Chapter 1) with the observation that the gap between words and the thoughts they are used to convey is far greater (and far more pervasive) than usually assumed: beyond cases of disambiguation ('I visited the bank today'), or reference assignment ('He's smart'), it also includes completion of discourse-initial subsentential utterances ('Bad day'), and various types of enrichment of the encoded content of the utterance ('I've had breakfast <today>', 'She's too young <for this job>', 'The park is some <considerable> distance from here'). In all of these instances, the meaning encoded by the linguistic expressions underdetermines the proposition expressed by the utterance ('what is said'). In order to arrive at what the speaker intended to say, hearers have to go beyond what is encoded by the linguistic string and augment, restrict or otherwise modify linguistically encoded meaning through a variety of inferential steps. It follows, then, that pragmatic processing is involved in recovering not only what the speaker is implicating, but also what proposition she is directly expressing.

Carston goes on to suggest that this discrepancy between what is linguistically encoded and what is said is a fundamental property of natural language. She argues convincingly against various 'effability' claims endorsed, among others, by Frege, Quine, Searle and Katz, according to which it is possible to construct sentences which encode a thought/proposition that is constant across all contexts: since pretty

0024-3841/03/$ - see front matter
doi:10.1016/S0024-3841(03)00035-4

much all natural-language sentences depend on contextual completion to express a complete proposition, the prospect of having (at least some) thoughts being fully encoded by linguistic strings turns out to be implausible. Further, Carston argues, utterances are usually interpreted against a massive set of weakly manifest (taken-for-granted) unrepresented assumptions and practices (what John Searle in his work calls the *Background)*. For instance, *open* receives different interpretations in 'John opened the can/the book/his mouth' according to different scenarios associated with these activities in our tacit world knowledge. The role of this set of background assumptions in utterance comprehension makes it even less likely that linguistic stimuli can be made to fully encode the thoughts they can be used to convey.

If linguistic underdeterminacy is a design feature of human communication, how is it, ultimately, resolved? Carston's answer is situated within the relevance-theoretic framework of communication (see Sperber and Wilson, 1986/1995). One of the main assumptions of the theory which lies at the heart of Carston's book is that human communication is first and foremost a kind of constrained mind-reading, an effort to reconstruct the speaker's intentions, with the linguistic code offering only partial and incomplete evidence for the thoughts being communicated. The gap between what is linguistically encoded and what is communicated is resolved through a pragmatic-inferential mechanism which invests processing effort in directions which hold the promise of relevance, or cognitive gains. In processing communicated stimuli, this pragmatic mechanism aims towards an optimal trade-off between processing effort and emerging cognitive effects: it therefore looks for cognitive effects in order of accessibility and stops when expectations of relevance have been satisfied.

With this background in place, Carston starts exploring the consequences for the architecture of the semantics-pragmatics interface (Chapter 2). On this view of linguistic communication, the linguistic meaning of an utterance (i.e. the output of the decoding process) typically consists of a set of incomplete logical forms (e.g. forms which contain slots for indexicals, unarticulated arguments, etc.). Such forms are the input to the pragmatic interpretation device which uses them as templates for the construction of fully propositional (complete, truth-evaluable) representations. The inferential development of full propositional forms out of these incomplete conceptual schemata is constrained by expectations of cognitive effects. To the extent that they fall within the speaker's communicative intentions, the assumption(s) that arise out of the development of the logical form of the utterance are *explicatures*. For instance, the explicature of (1) can be rendered as in (2):

(1) Spokesperson: He's here.
(2) President Bush is at the White House.

Explicatures can interact with contextual assumptions in order to produce *implicatures*. For instance, (2) may combine with the contextual premise in (3) to produce the implicature in (4):

(3) If President Bush is at the White House, the press conference can start.
(4) The press conference can start.

Both explicatures and implicatures, then, belong to the communicated content of an utterance. Their main difference lies in how this content is derived: explicatures are a product of both linguistic decoding and pragmatic inference, while implicatures rely solely on inference (although their derivation is constrained by the encoded content of the utterance). It is worth stressing that, within relevance theory, a single utterance may have more than one explicature; for instance, the proposition expressed by the utterance in (1) (i.e. (2)) may be embedded in various speech-act or propositional-attitude descriptions, thus giving rise to several (higher-level) explicatures:

(5) The spokesperson believes that President Bush is at the White House.
(6) The spokesperson is glad that President Bush is at the White House.

Identifying the explicature(s) and—where appropriate—the implicature(s) communicated by an utterance is the main goals of the hearer during a conversational exchange.

Carston's proposal differs in a number of ways from other current views on the semantics–pragmatics interface, and she devotes a large part of Chapter 2 to a thorough discussion of potential objections and alternatives to the relevance-theoretic position. One source of disagreement has centered on the precise number and type of representations involved in moving from an incomplete logical form to what corresponds to the explicature level. Kent Bach, for instance, has proposed a three-way distinction among 'what is said'/'impliciture'/'implicature', where impliciture corresponds broadly to the relevance-theoretic level of explicature and 'what is said' is kept an entirely semantic notion (with minimal pragmatic intrusions, mainly saturation of indexicals). A minimalist, mostly semantic construal of 'what is said' may have some appeal among philosophers but does not fit within a cognitively oriented pragmatic theory: as shown by Carston, even the saturation of an indexical such as *here* requires quite global considerations of what the speaker *meant* (since *here* can refer to this room, this building, this city, etc.). Furthermore, a narrowly defined notion of 'what is said' cannot deal with sentence fragments ('Typical', 'Telephone' etc.), since their constituents fall far short of expressing even a minimal proposition. For purposes of linguistic pragmatics, then, only two levels play a role in utterance comprehension: the logical form of the utterance on the one hand, and various communicated propositions (relevance-theoretic *explicatures* or *implicatures*) on the other.

A second source of controversy has been the type of pragmatic process involved in expanding the content of logical form into a complete, truth-evaluable propositional representation. Carston, alongside Sperber and Wilson, Bach, Recanati and several others, includes in such processes *disambiguation* (for structurally or lexically

ambiguous strings), *saturation* (for indexicals and other terms which require contextual filling-in of a slot in their semantic representation) and *free enrichment*, or the intrusion of contextual material even when no slot exists in the semantic representation. Other investigators, however, have recently argued that all cases of free enrichment should be reanalyzed as cases of saturation and all possible loci of pragmatic intrusion should be marked as such in the semantic entry of the corresponding lexical items (see Stanley, 2000, and for criticism, Bach, 2000). This view has immediate consequences for semantics, since it requires the presence of 'hidden variables' in the representation of a host of linguistic expressions (e.g. a comparison class variable for measure adjectives such as *long* or *short*, a domain variable for quantifiers such as *every* and *often*, etc.). Carston presents a range of compelling arguments against the 'no-free-enrichment' position, showing ultimately that, in a large number of cases, the presence of a hidden variable is implausible (e.g. 'Mary and John left <together>', 'Mary left John and <as a result> he became depressed'). Here, as in several other places throughout the book, Carston's conclusions have far-reaching implications for the task of semantic theory: since both free enrichment and saturation are pragmatic processes which operate on an utterance's schematic semantic output, the challenge for linguistic analysis is to decide which one is at work whenever some expression makes different contributions to truth conditions depending on context. Relevant examples include quantifier domain restriction ('Everyone <in the room> left'), the interpretation of the possessive ('John's novel'), or the temporal reference of the past perfect ('I've slept <this afternoon>'), where a consensus has yet to be reached on whether pragmatic input saturates a semantic gap or narrows the interpretation of a semantically more general expression.

Finally, an important question raised by this work is how to decide, for any pragmatically supplied element of meaning, whether it forms part of an explicature or an implicature. This is an immensely intriguing problem, and pretheoretic intuitions about what the speaker intended to say are only of partial help in solving it. After considering a variety of criteria, and discussing relevant proposals by Recanati and others, Carston concludes that the most useful tool for drawing the distinction is the Scope Principle: if an element of meaning is shown to fall under the scope of propositional connectives, it cannot be part of an implicature but has to contribute to the truth-conditional content of the utterance. This test shows, for instance, that quantifier domain restriction contributes to the proposition expressed by the utterance, and so does the temporal inference from the use of the perfect in (8):

(7) If everyone <in the room> left, then we can leave too.
(8) If you haven't slept <today>, you need some coffee.

The next three chapters offer a detailed illustration of the consequences of the underspecification thesis for the semantic and pragmatic analysis of specific phenomena: conjunction, negation, and on-line concept construction. Carston's discussion is

thorough and detailed; given that all of these areas have been the source of major controversies and intense theoretical interest within linguistics and philosophy, it is remarkable how the book manages both to do justice to the complexity of the issues and to navigate through the data with a solid sense of direction. Chapter 3 focuses on a range of well-known interpretations of the *and*-conjunction with special emphasis on the temporal and consequence readings (cf. 'We spent the day in New York and <then> I drove back', 'She shot him and <as a result> he died'). Carston makes a convincing case for not treating this systematic multiplicity of meanings as a case of semantic ambiguity; further, she uses the Scope Principle to argue that this is not a case of implicature either. Her main thesis is that *and* has its truth-functional semantics ('&'), and that the temporal and causal interpretations are derived through pragmatic enrichment which exploits (a) the script-like organization of our encyclopedic knowledge (e.g. assumptions about how events are normally ordered), and (b) our cognitive quest for causes and explanations. In Chapter 4, Carston tackles the complex problems posed by the various interpretive possibilities for negation in English, focusing on its scope properties and its descriptive vs. metalinguistic uses. This area has generated an immense literature and Carston discusses a lot of it, especially the work of Horn, Atlas, Grice, Gazdar, and Burton-Roberts. The logic of her main argument is similar to that of the previous chapter: she rejects the view that negation is semantically ambiguous or that narrow-scope readings are contributed by an implicature; she argues instead that negation has a wide-scope semantics which may be pragmatically narrowed in certain contexts. One of the most interesting parts of the discussion concerns metalinguistic negation. Carston proposes that negation preserves its truth-functionality even when metarepresentational (or 'echoic') material appears in its scope (cf. 'I was not happy; I was ecstatic'). Since this proposal first appeared (Carston, 1996), it has been very influential in discussions of metalinguistic negation. A minor complaint about the exposition here is that the author chooses not to repeat arguments she has previously published on the echoic nature of metalinguistic negation; nevertheless, it would be useful for the reader to have other relevant points collected in a single context.

The last chapter outlines a novel perspective on so-called 'loose use' of concepts. This covers cases where pragmatic processing does not narrow/enrich the propositional content of the utterance but loosens/broadens aspects of encoded conceptual content, as in metaphor ('He's a bulldozer') and hyperbole ('This steak is raw'). Carston departs from the standard relevance-theoretic position, according to which the effects of loosening are registered on the level of implicature alone. She suggests that concept loosening, just like concept narrowing, affects the content of the proposition expressed by the utterance; both processes offer ways of constructing on-line a concept which differs from the concept encoded by a word. Ultimately, of course, almost all lexically communicated concepts are departures from their encoded content in some way or another. On this radical view, lexically encoded concepts are schematic representations, or mere pointers to communicated concepts (see Sperber and Wilson, 2000, for discussion).

As is clear from this brief overview, *Thoughts and Utterances* addresses some of the deepest and most fundamental issues in the semantics–pragmatics interface and

the philosophy of language. The book brings together a wealth of empirical observations and new analyses and is impressive in breadth and depth. It is also one of the most detailed and powerful expositions of relevance theory and it enriches the framework in considerable ways. For these reasons, the book will be invaluable for both those researchers who work within relevance theory and those who approach problems in the semantics–pragmatics interface from different theoretical backgrounds.

One of the main contributions of Carston's specific proposals—and of the relevance-theoretic framework in general—is a rethinking of the borderline between semantics and pragmatics. As mentioned already, semantics in this framework is identified with a subcomponent of the grammar which delivers incomplete (typically sub-propositional) representations which form the input to processes of pragmatic interpretation. It follows that pragmatics, on this view, has an important role to play in the construction of the proposition expressed by the utterance—hence to the utterance's truth-conditional content. It also follows that the explicit/implicit distinction makes a significantly different cut than the semantics/pragmatics distinction (since pragmatic enrichments, loosenings and other post-semantic processes end up contributing to explicatures). The force and implications of these conclusions challenge traditional views of the semantics–pragmatics interface and bring communicative inference to the center of the study of linguistic meaning. This is an important shift: generally, even though some version of the semantic underdeterminacy thesis which recognizes the role of pragmatics in determining truth-conditional content has been accepted by many researchers (see, e.g., the work of Bach, S. Davis, Recanati, and Travis), its implications have not penetrated mainstream linguistic thought. As Carston shows, an approach which allows pragmatically inferred material to contribute to the proposition expressed can naturally accommodate a range of otherwise puzzling facts about utterance interpretation; for instance, it can explain the finding that post-decoding nonsemantic material can appear in the scope of logical operators (cf. the Scope Principle). More traditional approaches which limit all post-decoding nonsemantic material to the level of implicatures (cf. Levinson, 2000) are left with the uncomfortable conclusion that implicatures have to be allowed occasionally to have truth-conditional import. Hopefully, this book will lead to a wider appreciation of contextualist intrusions into 'what is said' by both semanticists and pragmaticists.

Most of the issues discussed in *Thoughts and Utterances* have traditionally been the topic of the philosophy of language; however, Carston's stance seeks to integrate the study of linguistic pragmatics with the broader program of contemporary cognitive science. Utterance interpretation is approached as a mental computational procedure defined over underlying semantic and post-semantic representations. Despite the fact that the book does not make much direct contact with empirical results from fields such as psycholinguistics, both Carston's main thesis and some of her more specific arguments can be shown to have broader consequences for the psychology of human communication.

First, the linguistic underdeterminacy thesis has some implications for the 'Whorfian turn' in modern cognitive science, which marks an interest in possible effects of systematic linguistic encoding on habitual thought patterns in humans.

Several of these discussions of language–thought relationships tacitly presuppose the existence of an isomorphism between linguistic strings and the thought contents they are used to express: they therefore assume that whatever is 'in the words' is shaped by (and, in turn, has the potential of shaping) whatever is 'in the thought'. In light of the fact that hearers routinely and massively use inference to go beyond the linguistic meaning of words and structures, this strong Whorfian possibility seems unlikely. For the same reasons, the underdeterminacy thesis is incompatible with the position, still popular in some philosophical circles, that the format of our thought representations (or Mentalese) just IS the natural language we happen to speak.

Another area which can be deeply affected by the proper treatment of linguistic underdeterminacy is language acquisition. One of the fundamental assumptions made by those who study word learning is that children come to grasp word meanings through systematic cross-situational observation (see, e.g., Pinker, 1984; for discussion, see Gleitman, 1990). However, the problem of discovering word meanings by mapping words onto real-world circumstances becomes much more complicated once it is acknowledged that the interpretation of words is not stable across contexts and that—as Carston further suggests—what is semantically encoded by a lexical item may be highly schematic and abstract. How can children isolate what is stable across the different learning opportunities in which a word occurs? Part of the solution to this complex question has to come from a fact which is at the heart of Carston's book: since children hear words in communicative contexts, their interpretation of words must be constrained by considerations of the intentions, dispositions etc. of the speaker. Word learning for children then has to rely heavily on some form of mind-reading. The implications of this position have led to some very interesting findings in the study of early lexical development (for a survey, see Bloom 2000, and the recent special issue of *Mind and Language* 17/2000 on *Pragmatics and Cognitive Science*). Related findings have begun to emerge in work which investigates the systematic contribution of semantic and pragmatic factors in older children's interpretation of lexical items (cf. Noveck, 2001; Papafragou and Musolino, 2003).

Finally, several of the ideas explored in this book lend themselves to psycholinguistic testing. To mention one example: unlike Grice and many others, who have taken 'what is said' to be both logically and computationally prior to 'what is implicated', Carston (and relevance theory) argues that the computation of explicatures does not, in fact, precede that of implicatures: rather, both types of communicated assumption are calculated through a process of 'mutual parallel adjustment' until the system's current expectations of relevance are met. Preliminary support for this claim comes from psycholinguistic evidence which shows that indirect requests and several 'non-literal' uses of language are understood much faster than what would be expected if their literal meaning was always accessed first (Gibbs, 1994; Clark, 1996). This position is also compatible with evidence which points to the massively interactive, probabilistic nature of the on-line comprehension system in adults (Trueswell and Tanenhaus, 1994). It would be interesting to see how the exact mental operations which underlie the 'mutual adjustment' of implicatures and explicatures can lead to empirically testable predictions. There are

several other topics throughout Carston's book which would merit experimental attention: the schema summarizing various interpretation paths for negation scope assignment (Chapter 4) can be made to support empirical (and falsifiable) predictions; the on-line pragmatic process of concept narrowing and loosening is another such example.

In sum: One of the most obvious features of human communication is the fact that what speakers *say* typically falls far short of what they *mean*. A perhaps less obvious fact is that the meaning of the words speakers *utter* typically falls far short of what they intend to *say*. Carston's book offers a detailed exploration of this second kind of linguistic underdeterminacy and its consequences for the mechanics of utterance interpretation. For anyone with a serious interest in semantics, pragmatics and the philosophy of language, this is a thought-provoking, highly original and very important contribution to the study of linguistic communication.

References

Bach, K., 2000. Quantification, qualification, and context: a reply to Stanley and Szabo. Mind and Language 15, 262–283.

Bloom, P., 2000. How Children Learn the Meaning of Words. MIT Press, Cambridge, MA.

Carston, R., 1996. Metalinguistic negation and echoic use. Journal of Pragmatics 25, 309–330.

Clark, H., 1996. Arenas of Language Use. University of Chicago Press, Chicago.

Gibbs, R., 1994. The Poetics of Mind. Cambridge University Press, Cambridge.

Gleitman, L., 1990. The structural source of verb meanings. Language Acquisition 1, 3–55.

Levinson, S., 2000. Presumptive Meanings. MIT Press, Cambridge, MA.

Noveck, I., 2001. When children are more logical than adults: experimental investigations of scalar implicature. Cognition 78, 165–188.

Papafragou, A., Musolino, J., 2003. Scalar implicatures: experiments at the semantics–pragmatics interface. Cognition 86, 253–282.

Pinker, S., 1984. Language Learnability and Language Development. Harvard University Press, Cambridge, MA.

Sperber, D., Wilson, D., 1986/1995. Relevance: Communication and Cognition, second ed. Blackwell, Oxford.

Sperber, D., Wilson, D., 2000. The mapping between the mental and the public lexicon. In: Carruthers, P., Boucher, J. (Eds.), Language and Thought: Interdisciplinary Themes. Cambridge University Press, Cambridge, pp. 184–200.

Stanley, J., 2000. Context and logical form. Linguistics and Philosophy 23, 391–434.

Trueswell, J., Tanenhaus, M., 1994. Toward a lexicalist framework of constraint-based syntactic ambiguity resolution. In: Clifton, C., Frazier, L. (Eds.), Perspectives on Sentence Processing. Erlbaum, Hillsdale, NJ, pp. 155–179.

Anna Papafragou
University of Pennsylvania
Institute for Research in Cognitive Science
3401 Walnut Str., Suite 400A
Philadelphia, PA 19104, USA
E-mail address: anna4@linc.cis.upenn.edu