

further. What he offered instead were rules of thumb, or maxims, by which one could determine the effects of common purpose. But these rules of thumb are no more than that. They are proxies, promissory notes, for a theory of collective action that is yet to be developed.

S&W, however, spurn Grice's observation and decline to develop any notion of "matter in hand" or "common purpose." When they characterize *relevance*, it is always divorced from what the participants in a discourse are really doing. "For us," say S&W (p. 161), "the only purpose that a genuine communicator and a willing audience necessarily have in common is . . . to have the communicator's informative intention recognised by the audience." Perhaps. But even that intention cannot be recognized without seeing the potential common purposes to which the communicators action is to contribute. S&W's position is like claiming that the only purpose I need in stepping on the car's accelerator is to put more gasoline into the carburetor. I do feed the carburetor, but that hardly accounts for why I usually take that action. I do it to speed up the engine, to turn the wheels faster, to speed up the car, to get me to my destination quicker, and so on. The "matter in hand," the main purpose for an action, is often quite remote from its immediate effects. Likewise in communication. That was Grice's and Austin's shared insight. So what S&W leave us with is a peculiarly empty notion of relevance. It almost belies the title of the book.

As a result, S&W pass off onto "cognitive psychology," without further explication, what for many scholars are the central issues of pragmatics. Consider their idea that addressees take the decoded content of an utterance along with what is "mutually manifest," weigh its "contextual effects" against its "processing effort," and select the interpretation that is "optimally relevant - that is, "the first accessible interpretation consistent with the principle [of relevance]" (Precis, sect. 3.3, para. 6). But what exactly are "accessibility" and "processing effort"? Do we really *select* interpretations, as they presuppose? How are contextual effects weighed against processing effort? And so on. Paradoxically, to answer these questions, we would need just the notions of collective action and evolving purpose that S&W are unwilling to provide.

Notions like these are already under investigation in the literature on conversation and other types of discourse, but S&W pass them over. In the end, however, they will not be able to duck the question "relevant to what?"

Layers of communication. Many types of discourse have more than one distinct layer of action or communication (Bruce 1981; Clark 1987; Goffman 1974). S&W, however, presuppose that all communication is flat - that it has only one layer, one type of relevance. That, I suggest, leads to a misrepresentation of many important phenomena in communication.

Layering is easiest to recognize in fiction - novels, plays, short stories, jokes, films - though it is common elsewhere too. In Herman Melville's novel *Moby Dick*, the narrator Ishmael is telling some landmen, perhaps in a Nantucket tavern, about his whaling adventures. Call this domain 1. Yet everything in domain 1 - Ishmael, his audience, his adventures, and his narrative - are Melville's inventions, and in writing the novel, he "communicated" that narrative to us readers. Call this domain 2. (One can argue for another domain between 1 and 2, but it is not needed to make the point.) Likewise, in *Hamlet*, Hamlet talks with Ophelia in a fictional domain 1, but Shakespeare "communicates" that conversation to us playgoers in domain 2 (via actors in yet another domain pretending to be Hamlet and Ophelia). All fiction has at least two domains, two layers.

How does relevance theory apply? In domain 1, the theory might claim that, fictionally, when Ishmael says, "Call me Ishmael," he has "informative" and "communicative intentions" toward his Nantucket audience, and they in turn presume that his utterance is relevant to them. The theory might also apply when Hamlet tells Ophelia, "Get thee to a nunnery."

What about domain 2? When Melville writes for us, "Call me Ishmael," we aren't to call *him* Ishmael. Nor is Shakespeare asking us (or Ophelia or the actress playing Ophelia) to hie off to some "nunnery" - some brothel. Domain 2 is somehow very different from domain 1.

Relevance theory doesn't go beyond the surface layer. It has nothing to say about domain 2, even if we avail ourselves of S&W's notion of "interpretation" or "interpretive representation." Here is why. Melville and Shakespeare have intentions toward us, but these are not "informative" or "communicative intentions" - they do not constitute the Gricean speaker's meaning (Clark 1987) - and the principle of relevance does not apply. So when Melville and Shakespeare "communicate" with us, it is communication of a fundamentally different type. This has a surprising but demonstrable consequence. Even if relevance theory could explain how the Nantucket landmen understood Ishmael, and how Ophelia understood Hamlet, it would not explain how *we* do. Relevance theory simply does not apply to a great deal of our most cherished communication.

Relevance theory has a long way to go to become a full theory of communication and cognition. It cannot work, I suggest, without well-developed notions of collective action and layering. But can it accommodate these without being stretched beyond the breaking point? That may be the next test of the trans-Channel alliance.

ACKNOWLEDGMENTS

This work is supported in part by NSF Grant BNS 83-20284.1 thank Eve V. Clark, Florence Edwards, and Ellen A. Isaacs for comments.

The task of the speaker and the task of the hearer

Anne Cutler

MRC Applied Psychology Unit, 15 Chaucer Road, Cambridge CB2 2EK, England

These brief remarks will be addressed to Sperber & Wilson's (S&W's) view of verbal communication. First, S&W draw a distinction between two separate processes of comprehension: a decoding process and an inferential process. They are principally concerned with the operation of the latter; the former they dismiss as automatic and therefore "not so much a part of the comprehension process as something that precedes the real work of understanding" (p. 177). Second, they imply that the work of understanding (though "real") is less than the work of speaking; the brunt of the work in communication is borne by the speaker. "It is left to the communicator to make correct assumptions about the codes and contextual information that the audience will have. . . . The responsibility for avoiding misunderstandings also lies with the speaker" (p. 43). "If the speaker has done her job properly, the end of the utterance should confirm all the provisional choices . . . that have been made along the route" (p. 208).

This picture is distinctly unfair to the hard-pressed hearer. Hearers are presented with signals which are for the most part semantically, syntactically, lexically, and phonologically unpredictable; moreover, the signals arrive in a noisy channel and are frequently subject to considerable distortion and attenuation. Speakers, on the other hand, have in principle a free hand in what they choose to say and how they choose to say it. S&W's principle of relevance is based on the observation that speakers do not take advantage of this freedom; in contrast, they constrain their utterances quite severely in order to make life easier for hearers.

In fact, S&W have here revealed only the tip of an iceberg. Speakers construct their speech output so as to cater to listeners' needs in a far more detailed fashion than is captured by the

guarantee of relevance or by Grice's injunctions to speakers not to bore, puzzle, offend, or deceive audiences. In particular, there is abundant evidence that speakers adjust their output to assist the listener at those levels which S&W claim are the subject of "automatic" processing - even at the level of segment production, as the following examples will show.

On the one hand, consider the inhibition of certain phonological rules of elision and assimilation. The application of such rules can result in a distortion, in casual speech, of phonetic segments which would be clearly articulated in more formal speech (Cooper & Paccia-Cooper 1980; Kaisse 1985); for example, the sequence [tj] can become the affricated segment [ts]. This palatalisation rule can apply across word boundaries, as in "Meetcha after work?" Cooper and Paccia-Cooper investigated the applicability of such palatalisation as a function of the informativeness of words preceding and following the boundary. For example, they varied word frequency of occurrence, comparing relatively common words ("rode your horse"; "had utensils") with much less frequent ones ("goad your horse"; "had euglena"). Varying the frequency of the word preceding the boundary had no effect on the frequency of palatalisation across the boundary; but varying the frequency of the word after the boundary had a strong effect - palatalisation was used significantly less often before rare words. Cooper and Paccia-Cooper also looked at the effect of contrastively stressing each word; again, stressing preceding words had no significant effect on the applicability of palatalisation, but stressing following words almost completely inhibited it.

In other words, distorting the ends of words is something speakers are fairly happy to do; but they are reluctant to distort word beginnings if the words are either rare or contrastively stressed, that is, if their information value is high. The beginning is the most important part of a word for the listener - distortion of word onsets disrupts word recognition far more than distortion of later segments (Bagley 1900; Cole 1973; Marslen-Wilson & Welsh 1978). So the speakers in Cooper and Paccia-Cooper's studies were clearly making phonological choices in such a way as to minimise disruption to the listener.

The same kind of motivation can be discerned in a pattern observed by Cutler (1983) in the correction of slips of the tongue. Errors of lexical stress occur quite frequently - *synTAX* for *SYNtax*, *origin* for *Origin*. Mostly such errors remain uncorrected by the speaker. This should cause the hearer little problem, since prosodic stress plays no role in word recognition (Cutler 1986); the hearer will probably notice a mismatch between spoken form and canonical lexical form, but will be readily able to discount it (cf. p. 23). What does disrupt word recognition, though, is getting vowel quality wrong - substituting a full for a reduced vowel or vice versa (Bond & Small 1983). So it is not surprising to find that precisely those stress errors which result in a change of vowel quality are the stress errors most likely to be corrected. Thus *origin*, in which a full vowel in the initial syllable has been replaced by a reduced vowel, and a reduced vowel in the second syllable has been replaced by a full vowel, is much more likely to be corrected by the speaker than *synTAX*, in which both vowels are full in both target and error.

These segment-level instances of perceptually driven speaking are striking; but one could easily add many instances at the lexical level (such as the tendency for nonce formations not to distort the real words on which they are based; Cutler 1980) or at the prosodic level (such as the fact that the greater the semantic contrast between a lexical slip of the tongue and the intended word, the more likely it is that the speaker will draw hearers' attention to a correction of that slip by stressing it; Levelt & Cutler 1983). Seen in this light, speakers' attention to ensuring relevance is merely one end of a continuum of hearer-coddling; there is certainly nothing special about it, and nothing that makes attention to hearers' inference processes qualitatively different from attention to hearers' decoding processes.

Thus, there is a sense in which the task of the speaker extends beyond the translation of a message into a spoken output;

speakers take upon themselves some responsibility for ensuring that hearers successfully accomplish understanding. But they do this purely out of self-interest, to ensure that their message gets across, and they do it precisely because the task of the listener is intrinsically so much harder than the task of the speaker. Particularly, it is harder at exactly those levels which Sperber & Wilson dismiss as the province of automatic and reflex processing. At these levels speakers strive to ensure reception of their message by hearers. Decoding is part of the work of understanding too.

Relevance and mutual knowledge

Martin Davies

Philosophy Department, Birkbeck College, University of London, Malet Street, London WC1E 7HX, England

It is common for philosophers of language to abstract from considerations of context dependence. The resulting picture of language is an idealised one; but the assumption behind much work over the last twenty years is that the idealisation does no serious theoretical damage. This assumption is, in turn, nourished by the idea that context dependence in general can be treated on the model of the indexicality exhibited by the word "I". The meaning of "I" specifies a very simple rule for assigning a reference relative to a context. Give or take "I," "here," "now," and a few other expressions to be treated on the same lines, the linguistic meaning of a sentence will determine a proposition expressed and, in particular, will determine truth conditions.

One of the major virtues of Sperber & Wilson's (S&W's) *Relevance* is its stress upon the fact that the common picture is not just idealised but mythological; truth conditions are radically underdetermined by linguistic meaning.

Whether or not the bold claims for relevance theory can be sustained at every point, the treatment of metaphor in *Relevance* and the subsequent paper (Sperber & Wilson 1986b) is a real advance. It shows that the apparent dichotomy between speaker-meaning accounts of metaphor following Searle (1979) and "seeing as" accounts following Davidson (1978) is spurious, thus correcting an impression given, for example, by Davies (1983). And it gives some determinate theoretical substance to the suggestion of Blackburn (1984, pp. 171-79) that a metaphor is an "invitation to explore" a comparison or image.

The radical underdetermination of truth conditions by linguistic meaning is enough to show that something is seriously wrong with the code model of communication: Even in the case in which there is only a single determinate thought to be communicated, the content of the thought is not fully encoded in the sentence uttered. But there is something else wrong with the code model as a model of human communication - that is, communication amongst creatures for whom there is a difference between entertaining a proposition, or having a proposition presented as a candidate for belief, on the one hand, and actually going forward in judgement and believing the proposition, on the other.

If propositions really were encoded in sentences, and I were equipped with a mechanism for decoding, then upon hearing an utterance of a sentence I would be presented with a proposition as a candidate for belief. This would be similar to the way in which, if I have a perceptual experience of the world as being a certain way, then the proposition that the world is that way is presented as a candidate for belief. In the case of perception, if I take my experience at face value, then I believe that the world is indeed that way. What is more, I usually do take my experience at face value. I do not require a justification for doing that; rather, I should need a reason not to take it so (cf. pp. 257-58, n. 28). But still, the difference between perception and belief remains. In the case of communication on the code model, it