

structural representations as values and will be subject to the principle of relevance – hence likely to have values which vary across limits set by the requirement of minimal processing cost.

Each of these consequences is strongly confirmed. It is well known that the ascription of truth-theoretic content to sentences is not compositional as defined over linguistic expressions, but involves a pervasive dependence on contextual factors (Partee 1984). If we assume that the articulation of such content is not part of the speaker's language faculty but is provided only for LT propositions, then the dependence of truth-theoretic content on contextual parameters is predicted. (Ascribing content to expressions of LT is no less a problem, but it is one that has to be faced anyway.)

Prediction (4) is confirmed by such anaphoric processes as using pronouns. These display a heterogeneity which has to be stipulated on truth-theoretic assumptions of linguistic content, their values ranging over any represented information accessible at minimal processing cost to the hearer. On the assumption that truth-theoretic interpretation is fixed pragmatically, constrained by the principle of relevance, we can provide a unitary account of the relatively weak linguistic content intrinsic to anaphoric expressions and we can directly predict the range of truth-theoretic values they display (Kempson 1986).

Predictions (2) and (1) jointly provide a rich vein of auxiliary predictions and confirmation. Recent studies in phonology, morphology, and syntax have provided increasing evidence that if a unitary explanation of the interpretation of natural language structures is to be given, it must assume at least one level of structure distinct from PF structure, s-structure, and LF (as defined in government/binding theory; Chomsky 1981). Gussenhoven (1983) has provided evidence that, though stress assignment interacts with hierarchical structure to determine interpretation, the level of structure required is some "semantic" one involving concepts such as predicate, argument, and focus domain, rather than that provided by syntax. In morphology there are many cases across widely varying languages in which the morphological structure does not provide the required basis for interpretation of the structure (Pesetsky 1983; Sadock 1985; Williams 1981). A simple example is *generative grammarian*, which requires the compounding of *generative* and *grammar* prior to interpreting the affix *-ian*, thus assuming for purposes of interpretation the structure: [[generative grammar] *-ian*].

And there are the problematic "Reconstruction" phenomena in which the Binding Principles (Chomsky 1981) constraining the interpretation of reflexives and pronouns apply not to LF structures but to the output of an "interpretation" process which in effect reverses Move WH (Williams 1986) – exemplified by *Which of each other's pictures do you think that Bill and Mary deliberately ruined?* None of these problems has been seen as related to the others, and the solutions proposed postulate various arrays of double-structure assignment. Yet each of the phenomena is predicted from Relevance assumptions, and for the same reason. In order for the claim that grammars are input systems to a central LT to be nontrivial, the mapping from phonological/morphological/syntactic structure onto logical form must carry the natural language structure onto a discrete structure in at least some cases. The cases which have emerged in these different areas are those in which the LT structure specified as the *logical form of the expression is not identical to that provided by the natural language structure*. The prediction of their existence is thus an immediate consequence of viewing grammars as input systems to a central device with its own LT (Kempson, forthcoming b). Furthermore, RT predicts the property of Reconstruction (embarrassing to GB) of not being subject to well-formedness conditions imposed on Move α . A process-mapping natural-language structures nontrivially onto structures of a discrete LT is by definition not restricted to the well-formedness conditions associated with processes internal to natural language.

The explanatory power of RT also extends to the functional motivation for properties of UG. Properties of UG such as the principle of Full Interpretation (that no expression at LF is contentless – Chomsky 1986) and the Theta criterion (providing a unique association between syntactic and logical arguments) can be seen as motivated by the constraint of maximizing relevance (Kempson, forthcoming a). This is not to deny their status as properties of UG, but rather to provide an explanation of their existence, the language faculty being hardwired to provide analyses of linguistic strings with no wasted processing effort.

Thus RT provides new and precise solutions to a whole range of old problems. In each case these solutions are not piecemeal devices but a coherent part of an overall explanation. There is now the need for a new linguistic theory articulating the mapping from phonological/morphological forms via syntactic structure onto LT expressions in order to give these putative solutions detailed empirical substance. But the general message is clear enough: RT opens up new avenues of research not merely in its own domains of philosophy of language, pragmatics, and psychology, but in semantics, syntax, and even morphology and phonology. In so doing, it proves itself to be a theory which offers considerable intellectual rewards.

Implicature explicated?

Stephen C. Levinson

Department of Linguistics, University of Cambridge, Cambridge CB3 9DA, England; Department of Linguistics, Stanford University, Stanford, Calif. 94305

Sperber & Wilson (S&W) claim that a single principle of relevance (R) can account for all pragmatic inference; but they fail to define R clearly enough to make this an empirical claim. R is a function of (measuring) contextual effects (E) balanced by processing costs (C) – related, say, as $R = E/C$ (though not thus subjectively assessed).¹ This allegedly yields determinate interpretations. Yet how the equation is solved varies: Sometimes R has a predetermined value (V), requiring contexts to be expanded until V is satisfied (p. 142); sometimes R is a comparative measure, selecting the best interpretation (pp. 144, 153); sometimes C has a threshold value such that the first accessible E-yielding context is automatically selected (p. 178). But if only R is fixed there is no determinate solution; if C is fixed too, then comparing interpretations may exhaust C without producing any solution. (A puzzle: If $C = n$ and $E = m$ for interpretation 1, but $C = n + 1$ and $E = m + 1$ for interpretation 2, then the cost of each will be the cost of comparing them, viz. $2n + 1$, leading to the conclusion that one should always pick maximum E regardless of C.)² Finally, without an objective measure of C, how can R be empirically tested?

There are further unclaritys: Is R a measure applied to surface structures (p. 204), semantic representations (SRs), or explicated logical forms (LFs) – or even recursively to LFs plus implicatures, or to all of these? If the answer is to more than one of these, then the claim that a speaker should optimize R makes no clear prediction about how the speaker should distribute "meaning" between the "encoded" and "implicated," and thus how a recipient can best interpret the utterance.

Turning to implicature: S&W's prototype cases are those corresponding to Grice's (1975a) relevance implicatures. S&W admit that their cases mostly involve implicated premises; they also assert (p. 83) that deduction is the key to understanding nondemonstrative inference, including implicature. But because implicated premises cannot be deduced, the very prototype relevance-implicatures cannot be accounted for. Nor will "accessibility" provide heuristics (see p. 201) in cases in which the implicated premise is, as often happens, not a recipient's (e.g. A's) assumption (cf. p. 194):

- A: "Would you drive a Zordia?"
 B: "I wouldn't drive ANY expensive car."
 Implicated premise: "A Zordia is an expensive car."

This failure is played down by introducing implicated premises as if they were contextual assumptions (e.g., counterfactual premise [5]e, p. 122). "Real" (pretheoretical) relevance is largely about guessing the speaker's goals (see Allen 1983; Dascal 1977; Holdcroft 1987), not maximizing information for R.

Turning to explicatures: These include all those generalized conversational implicatures (GCIs) whose understanding is the signal achievement of neo-Gricean pragmatics. GCIs are default pragmatic inferences which may be canceled by specific assumptions, but otherwise go through. S&W hardly engage this work (except on p. 262, note 4); yet they imply that R automatically accounts for GCIs as explicatures. The onus is on them of course. S&W consider that the neo-Gricean theory of GCIs (Atlas & Levinson 1981; Gazdar 1979; Horn 1972, 1984; Levinson 1983) is a result of brainwashing by the "code-model" of communication (pp. 36ff). But they make an important mistake in equating default but defeasible inferences with automatic decoding. Indeed, any theory of inference that takes processing cost into account must surely employ default inferencing – witness its almost universal use in computer models of language use.

The theory of GCIs makes clear predictions. If R also made clear predictions, they would almost certainly be different. The following may be a case:

- A: "If it's possible that the spy has *more* than two passports, he may yet escape."
 B: "He has two passports."

GCI theory holds that *two* means "at least 2" (allowing the possibility of more) but also has GCI "at most 2"; the theory thus predicts that B implicates "He has only two passports," thus suggesting that A's fears are groundless, and (perhaps) that the spy will not escape. If "two" means "at least two," R theory would presumably predict that there would be no explicature "at most two" because that would rule out the low-cost contextual implication "he may escape," thereby lowering R. Thus B's utterance should implicate "He may escape." My intuitions favour the GCI predictions.

S&W may complain that this argument allows implicatures to effect explicatures, whereas in fact explicatures form a prior independent level (see Kempson & Cormack 1981 and the effective rebuttal in Horn 1985). The explicature/implicature distinction is salutary as a means of emphasizing the pragmatic determination of LF, but it is not clear. The only criterion offered is that explicatures must contain the encoded SR or LF as a proper subpart (p. 181); but many implicatures meet that condition:

- A: "If Thatcher has won the election, she'll have won three times."
 B: "Thatcher has won."
 Implicature: "Thatcher has won three times."

Nor can explicature-calculation always precede implicature-calculation (cf. p. 179); for example, disambiguation must be guided by contextual effects.

S&W are to be congratulated for emphasizing the role of pragmatic inference in communication in general, and in propositional determination in particular. But other pragmaticists are unlikely to agree with monadic R theory (Sadock 1986). R is an informational measure; but Atlas and Levinson (1981), Horn (1984), and others have argued for the necessity of two counterbalancing informational measures. And, as noted, pretheoretical relevance is largely about satisfying the goals of others, whereas Quality (pace S&W) remains unreduced. For R to encompass all these other principles is to stretch it too thin.

Meanwhile, to compare R to, say, GCI theory, the newer formulation needs much clarification: Since R is a processing theory, how about a computer implementation?³

NOTES

1. This paragraph is not meant just to recycle the critique by Gazdar and Good (1982) (although to my mind S&W still haven't dealt fully with that), but rather to ask for clarification about how R can be computed.

2. S&W (p. 131) suggest that heuristics will allow estimation without computation. We need detailed hypotheses, however. Without them the puzzle looms large; with them, predictions might be made about specific kinds of misinterpretation.

3. One attempt at computer implementation is being made by R. Poznanski under K. Spark Jones of the Computer Laboratory, University of Cambridge.

The multidimensionality of pragmatics

James D. McCawley

Department of Linguistics, University of Chicago, Chicago, Ill. 60637

Sperber & Wilson (S&W) fairly quickly dismiss the notion of "mutual knowledge," for reasons that relate to problems created by the specific conception of mutual knowledge that they consider: one which is defined in terms of ordinary knowledge and thus requires reference to propositions of the form "A knows that B knows that A knows that B knows that . . ." ad infinitum.

Suppose, however, that one takes mutual knowledge to be something really mutual – that is, that one speaks of knowledge on the part of a group of two or more persons as a whole and does not attempt to reduce such knowledge to knowledge on the part of each of those persons individually. Mutual knowledge, thus understood, is not knowledge that the participants in a conversation hold individually about one another but rather knowledge they possess jointly. The identification of propositions as mutual knowledge does not, then, require the "infinite series of checks" that S&W speak of (p. 18). The following are unproblematic sources of mutual knowledge: (i) propositions that a participant has asserted during the conversation, unless a proposition has been either retracted by the speaker or challenged by another participant; (ii) information that is "manifest" (in the sense of S&W, p. 39) to the participants jointly, including information about what has been said so far in the conversation (e.g., the proposition that John said that pigs have wings is mutual knowledge, even though the proposition that pigs have wings isn't, if Mary challenged John's statement to that effect), and (iii) information commonly regarded as manifest to the members of a group in which the participants share membership. This last source of mutual knowledge is discussed insightfully by Prince (1978), who notes as an example that the identity of the mayor of Philadelphia is mutual knowledge if the participants in a conversation can reasonably assume that they all are residents of Philadelphia, as when strangers strike up a conversation on a Philadelphia bus. This last source of mutual knowledge is what licenses the uttering of (1) even if the speaker does not know whether the addressee has read *Sense and Sensibility*:

- (1) It's the sort of scene that would have made Marianne Dashwood swoon (*Relevance*, p. 44).

Sense and Sensibility has the status of a literary classic; and in circles in which familiarity with the classics is presumed, the identity and personality of their major characters is "manifest" and thus will count as mutual knowledge, in the sense that a person is entitled to utter (1) without first having to say who Marianne Dashwood is. Note that on this understanding, a proposition can be mutual knowledge in a given conversation even if not all of the participants know the proposition and even if the proposition is false: Many false propositions are manifest.

Despite S&W's dismissal of mutual knowledge, suggesting