Truthfulness and Relevance
Author(s): Deirdre Wilson and Dan Sperber
Source: *Mind*, New Series, Vol. 111, No. 443 (Jul., 2002), pp. 583-632
Published by: Oxford University Press on behalf of the Mind Association
Stable URL: http://www.jstor.org/stable/3093623
Accessed: 24/06/2014 20:14

Truthfulness and Relevance

DEIRDRE WILSON AND DAN SPERBER

This paper questions the widespread view that verbal communication is governed by a maxim, norm or convention of truthfulness which applies at the level of what is literally meant, or what is said. Pragmatic frameworks based on this view must explain the frequent occurrence and acceptability of loose and figurative uses of language. We argue against existing explanations of these phenomena and provide an alternative account, based on the assumption that verbal communication is governed not by expectations of truthfulness but by expectations of relevance, raised by literal, loose and figurative uses alike. Sample analyses are provided, and some consequences of this alternative account are explored. In particular, we argue that the notions of 'literal meaning' and 'what is said' play no useful theoretical role in the study of language use, and that the nature of explicit communication will have to be rethought.

1. Introduction

Here are a couple of apparent platitudes. As speakers, we expect what we say to be accepted as true. As hearers, we expect what is said to us to be true. If it were not for these expectations, if they were not often enough satisfied, there would be little point in communicating at all. David Lewis (who has proposed a convention of truthfulness) and Paul Grice (who has argued for maxims of truthfulness), among others, have explored some of the consequences of these apparent platitudes. We want to take a different line and argue that they are strictly speaking false. Of course hearers expect to be informed and not misled by what is communicated; but what is communicated is not the same as what is said. We will argue that language use is not governed by any convention or maxim of truthfulness in what is said. Whatever genuine facts such a convention or maxim was supposed to explain are better explained by assuming that communication is governed by a principle of relevance.

According to David Lewis (1975), there is a regularity (and a moral obligation) of truthfulness in linguistic behaviour. This is not a convention in Lewis's sense, since there is no alternative regularity which would be preferable as long as everyone conformed to it. However, for any language £ of a population *P*, Lewis argues that there is a convention of *truthfulness and trust in £* (an alternative being a convention of truthfulness and trust in some other language £′):

> My proposal is that the convention whereby a population *P* uses a language £ is a convention of *truthfulness* and *trust* in £. To be truthful in £ is to act in a certain way: to try never to utter any sentences of £ that are not true in £. Thus it is to avoid uttering any sentence of £ unless one believes it to be true in £. To be trusting in £ is to form beliefs in a certain way: to impute truthfulness in £ to others, and thus to tend to respond to another's utterance of any sentence of £ by coming to believe that the uttered sentence is true in £. (Lewis 1975, p. 167)

Lewis considers the objection that truthfulness might not be the only factor which needs to be taken into account, and replies as follows:

> *Objection*: Communication cannot be explained by conventions of truthfulness alone. If I utter a sentence σ of our language £, you—expecting me to be truthful in £—will conclude that I take σ to be true in £. If you think I am well informed, you will also conclude that probably σ is true in £. But you will draw other conclusions as well, based on your legitimate assumption that it is for some good reason that I chose to utter σ rather than remain silent, and rather than utter any of the other sentences of £ that I also take to be true in £. I can communicate all sorts of misinformation by exploiting your beliefs about my conversational purposes, without ever being untruthful in £. Communication depends on principles of helpfulness and relevance as well as truthfulness.
>
> *Reply*: All this does not conflict with anything I have said. We do conform to conversational regularities of helpfulness and relevance. But these regularities are not independent conventions of language; they result from our convention of truthfulness and trust in £ together with certain general facts—not dependent on any convention—about our conversational purposes and our beliefs about one another. Since they are by-products of a convention of truthfulness and trust, it is unnecessary to mention them separately in specifying the conditions under which a language is used by a population. (Lewis 1975, p. 185)

However, Lewis does not explain how regularities of relevance might be by-products of a convention of truthfulness. One of our aims will be to show that, on the contrary, expectations of truthfulness—to the extent that they exist—are a by-product of expectations of relevance.

Paul Grice (1967), in his *William James Lectures*, sketched a theory of utterance interpretation based on a Co-operative Principle and maxims of truthfulness, informativeness, relevance and clarity (Quality, Quantity, Relation and Manner). The Quality maxims went as follows:

(1) *Grice's maxims of Quality*

Supermaxim: Try to make your contribution one that is true.

(i) Do not say what you believe to be false. [*maxim of truthfulness*]

(ii) Do not say that for which you lack adequate evidence.

The supermaxim of Quality is concerned with the speaker's overall contribution (what is communicated, either explicitly or implicitly), while the first and second maxims of Quality relate only to what is said (the proposition explicitly expressed or asserted). Grice saw the first maxim of Quality, which we will call the maxim of truthfulness, as the most important of all the maxims. He says in the *William James Lectures*:

> It is obvious that the observance of some of these maxims is a matter of less urgency than is the observance of others; a man who has expressed himself with undue prolixity would, in general, be open to milder comment than would a man who has said something he believes to be false. Indeed, it might be felt that the importance of at least the first maxim of Quality is such that it should not be included in a scheme of the kind I am constructing; other maxims come into operation only on the assumption that this maxim of Quality is satisfied. While this may be correct, so far as the generation of implicatures is concerned, it seems to play a role not totally different from the other maxims, and it will be convenient, for the present at least, to treat it as a member of the list of maxims. (Grice 1967, p. 27)

In his 'Retrospective Epilogue', written 20 years later, this view is apparently maintained:

> The maxim of Quality, enjoining the provision of contributions which are genuine rather than spurious (truthful rather than mendacious), does not seem to be just one among a number of recipes for producing contributions; it seems rather to spell out the difference between something's being and (strictly speaking) failing to be, any kind of contribution at all. False information is not an inferior kind of information; it just is not information. (Grice 1989a, p. 371)

Notice, though, an interesting shift. While he talks of 'the maxim of Quality', Grice's concern here is with the speaker's contribution as a whole; indeed, there is room for doubt about whether he had the first maxim of Quality or the supermaxim in mind. We believe that this is not a minor detail. One of our aims is to show that the function Grice attributes to the Quality maxims—ensuring the quality of the speaker's overall contribution—can be more effectively achieved in a framework with no maxim of truthfulness at all.

There is a range of apparent counterexamples to the claim that speakers try to tell the truth. These include lies, jokes, fictions, metaphors and ironies. Lewis and Grice are well aware of these cases, and discuss them in some detail. Grice (1967, p. 30), for instance, notes that his maxims may be violated, and lists several categories of violation, each with its characteristic effects. Lies are examples of *covert violation*, where the hearer is meant to assume that the maxim of truthfulness is still in force and that the speaker believes what she has said. Jokes and fictions might be seen as cases in which the maxim of truthfulness is overtly *suspended* (the speaker overtly *opts out* of it); the hearer is meant to notice that it is no longer operative, and is not expected to assume that the speaker believes what she has said. Metaphor, irony and other tropes represent a third category: they are *overt violations* (*floutings*) of the maxim of truthfulness, in which the hearer is meant to assume that the maxim of truthfulness is no longer operative, but that the supermaxim of Quality remains in force, so that some true proposition is still conveyed.

We will grant that a reasonable—if not optimal—treatment of lies, jokes and fictions might be developed along these lines. Tropes, and more generally loose uses of language (e.g. approximations, sense extensions), present a much more pressing challenge. After all, many, if not most, of our serious declarative utterances are not strictly and literally true, either because they are figurative, or simply because we express ourselves loosely.

An utterance can be said to have a literal meaning which is capable of being either true or false when the result of combining its linguistic sense with its reference is a proposition. We ourselves do not claim that all utterances have a literal meaning, and we will be arguing that even when a literal meaning is available, it is not automatically the preferred interpretation of an utterance. In fact, literalness plays no role in our account of language comprehension, and nor does the notion of *what is said*. By contrast, to those who argue that there is an expectation of truthfulness in what is said, literal meanings matter. For Grice, what is said (as distinct from what is implicated) is the literal meaning of an utterance. For Lewis, what is said is specifiable on the basis of the utterance's literal meaning. Without such an appeal to literal meaning in the determination of what is said, the claim that there is a maxim or convention of truthfulness in what is said would be, if not vacuous, at least utterly vague.

2. The case of tropes

Lewis (1975) considers the case of tropes:

> *Objection*: Suppose the members of a population are untruthful in their language £ more often than not, not because they lie, but because they go in heavily for irony, metaphor, hyperbole, and such. It is hard to deny that the language £ is used by such a population.
>
> *Reply*: I claim that these people are truthful in their language £, though they are not *literally truthful* in £. To be literally truthful in £ is to be truthful in another language related to £, a language we can call literal-£. The relation between £ and literal-£ is as follows: a good way to describe £ is to start by specifying literal-£ and then to describe £ as obtained by certain systematic departures from literal-£. This two-stage specification of £ by way of literal-£ may turn out to be much simpler than any direct specification of £. (Lewis 1975, p. 183)

Lewis's reply rests on a widely-shared view which dates back to classical rhetoric. On this view,

(2) (a) Figurative and literal utterances differ not in the kind of meanings they have (thus, if literal meanings are truth-conditional, so are figurative meanings), but in the way these meanings are generated.

(b) The meanings of figurative utterances are generated by systematic departures from their literal meanings.

For example, consider (3) and (4), where (3) is a metaphor and (4) is intended as hyperbole:[1]

(3) The leaves danced in the breeze.

(4) You're a genius.

Lewis might want to say that in literal-English, sentences (3) and (4) have just their literal meanings. In actual English, the language in which a convention of truthfulness and trust holds among English speakers, (3) and (4) are ambiguous. They have their literal meanings plus other, figurative meanings: thus, (3) has the metaphorical meaning in (5), and (4) the hyperbolical meaning in (6):

(5) The leaves moved in the breeze as if they were dancing.

[1] In this paper we will focus on metaphor, hyperbole and a range of related phenomena. For analyses of irony and understatement, see Sperber and Wilson (1981; 1986a, Ch.4, Scs.7, 9; 1990; 1998b); Wilson and Sperber (1992).

(6) You're very clever.

However, it is not as if any language £ (in the sense required by Lewis, where the sentences of £ can be assigned truth-conditional meanings) had ever actually been specified on the basis of a corresponding 'literal-£'. So what is the justification for accepting something like (2a) and (2b)? How are figurative meanings derived from literal meanings, and under what conditions do the derivations take place? Lewis does not explain, and there are no generally accepted answers to these questions. We have argued (Sperber and Wilson 1986a,b; 1990; 1998a) that figurative interpretations are radically context-dependent, and that the context is not fixed independently of the utterance but constructed as an integral part of the comprehension process. If so, then the very idea of generating the sentences of a language £ on the basis of a corresponding 'literal-£' is misguided.

Grice is often seen as providing an explanation of how figurative interpretations are conveyed. Consider a situation where the speaker of (3) or (4) manifestly could not have intended to commit herself to the truth of the propositions literally expressed: it is common knowledge that she knows that leaves never dance, or that she does not regard the hearer as a genius. She is therefore overtly violating the maxim of truthfulness: in Grice's terms, she is *flouting* it. Flouting a maxim indicates a speaker's intention: the speaker intends the hearer to retrieve an implicature which brings the full interpretation of the utterance (i.e. what is said plus what is implicated) as close as possible to satisfying the Co-operative Principle and maxims. In the case of tropes, the required implicature is related to what is said in one of several possible ways, each characteristic of a different trope. With metaphor, the implicature is a simile based on what is said; with irony, it is the opposite of what is said; with hyperbole, it is a weaker proposition, and with understatement, a stronger one. Thus, Grice might analyse (3) as implicating (5) above, and (4) as implicating (6).

Note that this treatment of tropes does not differ radically from Lewis's, or from the classical rhetorical account. Grice's approach, like Lewis's, is based on assumption (2a) and, more importantly, on assumption (2b) (that the meanings of figurative utterances are generated by systematic departures from their literal meanings). The only difference is that Lewis sees these departures as systematic enough to be analysed in code-like terms: the figurative meaning of a sentence is a genuine linguistic meaning specified in the grammar of £ by a derivation which takes the literal meaning of the sentence as input. The sentences of £ (unlike those of literal-£) are systematically ambiguous between literal and figurative senses. For Grice, by contrast, sentences have only literal meanings. Figurative meanings are not sentence meanings but utterance meanings, derived in a conversational context. However, the derivations proposed in Grice's pragmatic approach to tropes are the same as those hinted at by Lewis in his linguistic approach, and neither differs seriously from the classical rhetorical account.

Grice's treatment of tropes leaves several questions unanswered, and we will argue that it is inconsistent with the rationale of his own enterprise. In particular, there is room for doubt about what he meant by the maxim of truthfulness, and the role it was intended to play in his framework. This doubt is created by two possible interpretations of his notion of *saying*. On the first interpretation, saying is merely expressing a proposition, without any necessary commitment to its truth. Understood in this way, the maxim of truthfulness means 'Do not express propositions you believe to be false.' The function of this maxim, and more generally of the Quality maxims, would be to account for the fact—to the extent that it is a fact—that a speaker actually commits herself to the truth of what she says. Tropes would then be explained by the claim that flouting the maxim triggers the recovery of an implicature in the standard Gricean way. However, there is a problem. In general, the recovery of implicatures is meant to restore the assumption that the maxims have been observed, or that their violation was justified in the circumstances (as when a speaker is justified by her ignorance in providing less information than required) (Grice 1989a, p. 370). In the case of tropes, the maxim of truthfulness is irretrievably violated, and the implicature provides no circumstantial justification whatsoever.

On the second, and stronger, interpretation, saying is not merely expressing a proposition but asserting it: that is, committing oneself to its truth. Understood in this way, the maxim of truthfulness means 'Do not *assert* propositions you believe to be false.' On this interpretation, saying already involves speaker commitment, and the function of the maxim of truthfulness, and more generally of the Quality maxims, would be to ensure that speakers do not make spurious commitments. This seems to fit with Grice's above remark that the function of the Quality maxim is to guarantee that contributions are genuine rather than spurious. However, understood in this way, it is hard to see why a maxim of truthfulness is needed at all. It seems to follow from the very notion of an assertion as a commitment to truth (perhaps together with a proper understanding of commitment) that your assertions should be truthful. In fact, the only pragmatic function of the maxim of truthfulness, on this interpretation, is to be violated in metaphor and irony, thus triggering the search for an implicature. Without it, Grice would have no account of figurative utterances at all.

Which notion of saying did Grice have in mind in proposing the maxim of truthfulness? There is evidence of some hesitation. On the one hand, he treats the tropes as 'Examples in which the first maxim of Quality is flouted' (Grice 1967, p. 34). On the other, he comments that in irony the speaker 'has said *or has made as if to say*' (our italics) something she does not believe, and that in metaphor what is communicated must be obviously related to what the speaker '*has made as if to say*' (ibid. p. 34). If the speaker of metaphor or irony merely 'makes as if to say' something, then the stronger notion of saying must be in force; on the other hand, if the speaker of a trope merely 'makes as if' to say something, then surely the maxim of truthfulness is not violated. But if the maxim of truthfulness is not violated, how does Grice's analysis of metaphor and irony go through at all?

Elsewhere in his philosophy of language, where the notion of saying plays a central role, it was the stronger rather than the weaker notion that interested Grice. He says, for example, 'I want to say that (1) "*U* (utterer) said that *p*" entails (2) "*U* did something *x* by which *U* meant that *p*".' (Grice 1967, p. 87). For Grice, what is *meant* is roughly coextensive with what is *intentionally communicated*: that is, with the information put forward as true. On this interpretation, saying involves speaker commitment: that is, it means asserting. Among his commentators, Stephen Neale (1992) treats these broader considerations as decisive: 'If *U* utters the sentence "Bill is an honest man" ironically, on Grice's account *U* will *not* have said that Bill is an honest man: *U* will have made as if to say that Bill is an honest man.' (Neale 1992, section 2).

How can we reconcile these two claims: that metaphor and irony are deliberate violations of the maxim of truthfulness, and hence must *say* something, and that in metaphor and irony the speaker merely *makes as if to say* something? A possible answer would be to distinguish two phases in the utterance interpretation process. In the first, the utterance of a declarative sentence would provide prima facie evidence for the assumption that an assertion is being made. In the second, this assumption would be evaluated and accepted or rejected. In the case of metaphor and irony, this second phase would involve an argument of the following sort: if it is common ground that the utterer *U* doesn't believe *p*, then *U* cannot assert *p*; it is common ground that *U* doesn't believe *p*; hence, *U* hasn't asserted *p*. In this way, we get a consistent interpretation of the notion of saying, and we can see why Grice hesitates between 'saying' and 'making as if to say.'

However, if this interpretation is correct, then a trope involves no real violation of the maxim of truthfulness at any stage: since the speaker was not saying *p*, she was not saying what she believed to be false. A flouting, so understood, is a mere appearance of violation. So why should it be necessary to retrieve an implicature in order to preserve the assumption that the maxims have been respected? The Gricean way to go (although Grice himself did not take this route) would be to argue that it is not the maxim of truthfulness but some other maxim that is being violated. Quite plausibly, the maxim of Relation ('Be relevant') is being violated, for how can you be relevant when you speak and say nothing? Surely the first maxim of Quantity ('Make your contribution as informative as is required') is being violated, for if nothing is said, no information is provided. The implicature should thus be seen as a way of providing a full interpretation of the utterance in which these maxims are respected.

The problem with this analysis of tropes (and with the alternative analysis on which floutings of the maxim of truthfulness are genuine violations) is that it leads to an interpretation of figurative utterances which irretrievably violates the Manner maxims. In classical rhetoric, a metaphor such as (3) or a hyperbole such as (4) is merely an indirect and decorative way of communicating the propositions in (5) or (6). This ornamental value might be seen as explaining the use of tropes, insofar as classical rhetoricians were interested in explanation at all. Quite sensibly, Grice does not appeal to ornamental value. His supermaxim of Manner is not 'Be fancy' but 'Be perspicuous.' He assumes, this time in accordance with classical rhetoric, that figurative meanings, like literal meanings, are fully propositional, and always paraphrasable by means of a literal utterance. Which raises the following question: isn't a direct and literal expression of what you mean *always* more perspicuous (and in particular less obscure and less ambiguous, cf. the first and second Manner maxims) than an indirect figurative expression? (Remember: you cannot appeal to the subtle extra effects of tropes, since they are not considered, let alone explained, within the Gricean framework.)

It would be presumptuous to attribute Grice's apparent hesitation between two senses of *saying* to a lack of conceptual rigour on his part. We see it rather as arising from the difficulty of deploying a notion of *saying* which is both close enough to common usage to justify the choice of this word, and yet precise enough to make a contribution to a theory of language use. We will argue in §7 that there is no such notion. We ourselves do not use 'saying' as a theoretical term, except in rendering the views of others.

3. The case of loose use

Tropes are the most striking examples of serious utterances where the speaker is manifestly not telling the strict and literal truth. Even more common are loose uses of language (e.g. rough approximations, sense extensions), where an expression is applied to items that fall outside its linguistically-determined denotation, strictly understood. Consider the examples in (7)–(10):

(7) The lecture starts *at five o'clock.*

(8) Holland is *flat.*

(9) *Sue:* I must *run* to the bank before it closes.

(10) *Jane:* I have a terrible cold. I need a *Kleenex.*

If the italicized expressions in (7)–(10) are understood in the most restrictive way (and ignoring issues of ambiguity or polysemy for a moment), these utterances are not strictly and literally true: lectures rarely start at exactly the appointed time, Holland is not a plane surface, Sue must hurry to the bank but not necessarily run there, and other brands of disposable tissue would do just as well for Jane. Such loose uses of language are very common. Some are tied to a particular situation, produced once and then forgotten. Others may be regular and frequent enough to give rise to an extra sense, which may stabilize in an individual or a population: lexical broadening (along with lexical narrowing and metaphorical transfer) has been seen as one of the main pragmatic factors driving semantic change (Lyons 1977, Chs 13.4, 14.5). What concerns us here is not so much the outcome of these historical macro-processes as the nature of the pragmatic micro-processes that underlie them, and we will largely abstract away from the question of whether, or when, a word such as 'flat', or 'run', or 'Kleenex' may be said to have acquired an extra stable sense (see Sperber and Wilson 1998a for some discussion).

How should loose uses such as those in (7)–(10) be analysed? Are they lapses, the result of sloppy speech or thought, accepted by hearers whose expectations have been reduced to realistic levels by repeated encounters with normal human failings? Is it reasonable to assume that there really is a convention or maxim of truthfulness, although speakers quite commonly fall short of strictly obeying it? As hearers, would we always—and as speakers, should we always—prefer the strictly true statements in (11)–(14) to the loose uses in (7)–(10)?

(11) The lecture starts at or shortly after five.

(12) Holland has no mountain and very few hills.

(13) I must go to the bank as fast as if I were running.

(14) I need a Kleenex or other disposable tissue.

Clearly not. In most circumstances, the hearer would not be misled by strictly untrue approximations such as (7)–(10), and their strictly true counterparts in (11)–(14) would not provide him with any more valuable information. Indeed, since these strictly true counterparts are typically longer, the shorter approximations may be preferable.

Loose uses of language present few problems for speakers and hearers, who are rarely even aware of their occurrence; but they do raise a serious issue for any philosophy of language based on a maxim or convention of truthfulness. We have suggested above that appeals to ambiguity (or polysemy) merely defer the problem, since such ambiguities ultimately derive from repeated instances of loose use (for further discussion, see §6 below). In this section, we will consider other solutions proposed in the literature, paying particular attention to Lewis's treatment of pragmatic vagueness. We will argue that no single solution, nor any combination of proposed solutions, is adequate to handle the full variety of loose uses of language, which go well beyond the types of pragmatic vagueness dealt with on Lewis's account.

In Grice's framework, loose uses such as (7)–(10) apparently violate either the maxim of truthfulness or the second maxim of Quality ('Have adequate evidence for what you say'). However, they do not really fit into any of the categories of violation listed in §1 above. They are not covert violations, designed to deceive the hearer into believing the proposition strictly and literally expressed. They are not like jokes or fictions, which suspend the maxims entirely. One might try to analyse them as floutings: overt violations (real or apparent), designed to trigger the search for a related implicature (here a hedged version of what was literally said or quasi-said); but the problem is that loose uses are not generally perceived as violating the Quality maxims at all. In classical rhetoric, they were not treated as tropes involving the substitution of a figurative for a literal meaning. They do not have the striking quality that Grice associated with floutings, which he saw as resulting in figurative or quasi-figurative interpretations. Loose talk involves no *overt* violation, real or apparent; or at least it does not involve a degree of overtness in real or apparent violation which might trigger the search for an implicature. While we are all capable of realizing on reflection that utterances such as (7)–(10) are not strictly and literally true, these departures from truthfulness pass unattended and undetected in the normal flow of discourse. Grice's framework thus leaves them unexplained.

Perhaps we should reconsider the apparent platitudes we started with. Maybe we should have said that as speakers, we expect what we say to be accepted as *approximately* true, and as hearers, we expect what is said to us to be *approximately* true. But this is far too vague to do the required work of explaining how speakers and hearers manage to communicate successfully. Approximations differ both in kind and in degree, and their acceptability varies with content and context. There is no single scale on which the degrees of approximation in disparate statements such as (7)–(10) can be usefully compared. The same statement can be an acceptable approximation in one situation and not in another. Thus, suppose the speaker of (7) expects the lecture to start sometime between five o'clock and ten past five: then (7) would be an acceptable approximation to a student who has just asked whether the lecture starts at five or six o'clock, but not to a radio technician preparing to broadcast the lecture live. Moreover, as we will argue below, there are cases where the notion of 'degrees of approximation' does not really apply.

A convention of truthfulness and trust in a language (if there were one) might play a valuable role in explaining how linguistic expressions acquire their conventional meanings, and how speakers and hearers use these meanings to communicate successfully. If all that speakers and hearers are entitled to are vague expectations of approximate truth, it is hard to see how the resulting convention of approximate truthfulness could be robust enough to establish common meanings. As we have shown, the same approximation may be differently intended and understood in different circumstances. Unless it is supplemented by some account of how speakers and hearers may converge on these more specific understandings—an account which might then be doing most of the explanatory work—a convention of approximate truthfulness and trust is inadequate to explain how the co-ordination necessary for successful communication is achieved. Still, this is the direction that David Lewis proposes to explore. He writes:

> When is a sentence true enough? [...] this itself is a vague matter. More important for our present purposes, it is something that depends on context. What is true enough on one occasion is not true enough on another. The standards of precision in force are different from one conversation to another, and may change in the course of a single conversation. Austin's 'France is hexagonal' is a good example of a sentence that is true enough for many contexts but not true enough for many others. (1979, pp. 244–45)

We agree with Lewis (and Unger 1975, Ch.2) that 'hexagonal' and 'flat' are semantically absolute terms, and that their vagueness should be seen as pragmatic rather than semantic (so in our terms they are genuine cases of loose use). However, Lewis's analysis of pragmatically vague terms such as 'flat' is very similar to his analysis of semantically vague terms such as 'cool'.[2] For Lewis, a semantically vague term has a range of possible sharp delineations, marking different cut-off points between, say, 'cool' and 'warm'. 'This is cool' may be true at some but not all delineations, and depending on our purposes, we may be willing or unwilling to assert it: hence its vagueness (1970, pp. 228–29). On Lewis's account, semantically absolute but pragmatically vague terms are handled on similar lines, except that semantic delineations are replaced by contextually-determined standards of precision (so if 'flat' were semantically rather than pragmatically vague, the analysis would not be very different). On this approach, 'Holland is flat' would be true according to some fairly low standard of precision, but false given higher standards.

Semantic vagueness clearly exists ('bluish' and 'flattish' are good examples); its analysis raises problems of its own, about which we have nothing to say here (see Williamson 1994; Keefe and Smith 1996). What we do want to argue against is the idea that loose use can be successfully treated as a pragmatic analogue of semantic vagueness. As we have suggested above (and will argue in more detail below), there are many varieties of loose use, not all of which can be satisfactorily handled by appeal to contextually-determined standards of precision. For the cases that cannot be handled on Lewis's lines, an alternative analysis must be found. We will propose such an analysis, and argue that it generalizes straightforwardly to all varieties of loose use (and indeed to all utterances, literal, loose, or figurative), making the appeal to standards of precision as a component of conversational competence unnecessary.

In fact, there are problems even in some cases where the appeal to contextually-determined standards of precision looks initially plausible. Consider a situation where (7) ('The lecture starts at five o'clock')

[2] For discussion, see Gross (1998).

would be accepted as true enough if the lecture started somewhere between five o'clock and ten past five. On Lewis's account, it might be claimed that here the contextually-determined standard of precision allows for a give or take of, say, fifteen minutes around the stated time. It should then follow that a hearer in the same situation, with the same standard of precision in force, would be equally willing to accept (7) as true enough if the lecture started somewhere between ten to five and five o'clock. But there is an obvious asymmetry between the two cases. Intuitively, the reason is clear enough: the audience won't mind or feel misled if they get to the lecture a few minutes early, but they will if they get there a few minutes late, so the loosening is acceptable only in one direction. In a different situation—when the speaker is talking about the end of the lecture rather than the beginning, for example—there may be an asymmetry in the other direction. Again, the reason is intuitively clear: the audience won't mind or feel misled if they can get away a bit earlier than expected, but they will if they have to stay longer. It is hard to explain these obvious intuitions by appeal to the regular notion of contextually-determined standards of precision as described above. One might, of course, try building the asymmetries into the standards of precision themselves, but then two different standards would have to be invoked to explain how (15) is quite naturally understood to mean something like (16):

(15) The lecture starts at five o'clock and ends at seven o'clock.

(16) The lecture starts at five o'clock or shortly after and ends at seven o'clock or shortly before.

This is clearly *ad hoc*. It would be better to find an alternative account of these asymmetries—but such an account might make the appeal to contextually-determined standards of precision redundant.

A more serious problem for Lewis is that in some cases of loose use, the appeal to contextually-determined standards of precision does not seem to work at all. Lewis's account works best when there is a continuum (or ordered series) of cases between the strict truth and the broadest possible approximation. 'Flat' is a good example, since departures from strict flatness may vary in degree. 'Five o'clock' also works well in this respect, since departures from exactness may vary in degree. But with 'run' in (9) ('I must run to the bank') and 'Kleenex' in (10) ('I need a Kleenex'), no such continuum exists. There is a sharp discontinuity between running (where both feet leave the ground at each step) and walking (where there is always at least one foot on the ground). Typically (though not necessarily), running is faster than walking, so that 'run' may be loosely used, as in (9), to indicate the activity of going on foot (whether walking or running) at a speed more typical of running. But walking at different speeds is not equivalent to running relative to different standards of precision. Similarly, 'Kleenex' is a brand name: other brands of disposable tissue are not Kleenex. The word 'Kleenex' may be loosely used, as in (10), to indicate a range of tissues similar to Kleenex. But there is no continuum on which being similar enough to Kleenex amounts to actually being Kleenex relative to stricter or looser standards of precision. 'Run', 'Kleenex' and many other words have sharp conceptual boundaries and no ordered series of successively broader extensions which might be picked out by raising or lowering some standard of precision. Yet these terms are often loosely used. This supports our claim that looseness is a broader notion than pragmatic vagueness.[3]

Again, for someone with no particular theoretical axe to grind, it is easy enough to see intuitively what is going on. Suppose you have a lecture one afternoon, but don't know exactly when it is due to start. Someone tells you, 'The lecture starts at five o'clock.' From the literal content of the utterance, together with other premisses drawn from background knowledge, you can derive a number of conclusions that matter to you: that you will not be free to do other things between five and seven o'clock, that you should leave the library no later than a quarter to five, that it will be too late to go shopping after the lecture, and so on. To say that these conclusions matter to you is to say that you can use them to derive still further non-trivial contextual implications, of a practical or a theoretical nature. These initial conclusions are the main branches of a derivational tree with many further branches and sub-branches. You would have been able to derive all these direct and indirect conclusions from the strictly true utterance 'The lecture starts at or shortly after five o'clock,' but at the extra cost required to process a longer sentence and a more complex meaning. There are other conclusions—false ones this time—that you would have been able to derive from the approximation, 'The lecture starts at five o'clock', but not from its strictly true counterpart: that the lecture will have begun

[3] Since what we are calling 'loose uses' shade off into figurative uses such as hyperbole and metaphor, it might be argued that Lewis's analysis of vagueness could be saved by treating 'run' and certain other examples which present problems for his analysis as falling on the figurative side (cf. Gross 1998). But this would merely transfer them from one problematic category to another since, as we have argued, neither Lewis nor Grice has proposed a satisfactory analysis of tropes. Moreover, the move would be *ad hoc* since, as we have also shown, these uses have little in common with standard examples of metaphor and hyperbole recognized in the literature on rhetoric.

by one minute past five, for instance. But you are unlikely to derive them. They don't matter, because they are derivational bare branches which yield no further non-trivial implications.

Suppose that Peter and Mary, who are both rather unfit, are discussing where to go on their next cycling holiday. Mary suggests Holland, adding, 'Holland is flat.' From the strictly false proposition that Holland is flat—just as easily as from the strictly true hedged proposition that Holland is approximately flat—Peter can derive the true conclusion that cycling in Holland would involve no mountain roads and would not be too demanding. Unlike the true hedged proposition, the false approximation also has false implications (that there are no hills at all in Holland, for instance). But it is unlikely that Peter would even contemplate deriving any of these.

Suppose that Sue, chatting with friends in the street, looks at her watch and says, 'I must run to the bank before it closes.' Her friends will take her to mean that she must break off their chat and hurry to the bank. For them, that much information is worth deriving. Whether she will actually get to the bank by running, walking fast or a mixture of both is of no interest to them, and they will simply not attend to this aspect of the literal meaning of her utterance.

Suppose that Jane and Jack are in the cinema waiting for the film to start. By saying, 'I have a terrible cold. I need a Kleenex,' Jane provides a premiss from which Jack can infer that she wants to borrow a tissue to use in dealing with her cold. Her utterance also provides a premiss from which he could derive the possibly false conclusion that she does not want a tissue of any other brand than Kleenex; but he is unlikely to draw such a conclusion, since his expectations of relevance in this context are satisfied by the weaker interpretation on which she wants a tissue (and in a context where his expectations of relevance would encourage an interpretation on which Jane was specifically requesting a Kleenex—e.g. if she was angrily throwing away tissues of another brand—the utterance would not be understood as a case of loose talk).

As these examples show, hearers have no objection to strictly false approximations as long as the conclusions they bother to derive from them are true. In fact, they might prefer the shorter approximations to their longer-winded but strictly true counterparts for reasons of economy of effort. There is some evidence that speakers take account of the perceived preferences of their audience in deciding how strictly or loosely to speak. In a series of experiments on truthfulness and relevance in telling the time, Van der Henst, Carles and Sperber (forthcoming) showed that when people in public places are asked the time by a stranger, they tend to respond with a time that is either accurate to the minute or rounded to the nearest multiple of five, depending on how useful in the circumstances they think a more accurate answer would be. This includes people with digital watches, from which a time that is accurate to the minute is most easily read off. These people have to make an extra effort to produce a rounded answer which, to the best of their knowledge, is not strictly true, but is easier for their audience to process.

Anticipating the arguments of the next section, let us say that an utterance is *relevant* when the hearer, given his cognitive dispositions and the context, is likely to derive some genuine knowledge from it (we will shortly elaborate on this). Someone interested in defending a maxim or convention of truthfulness might then suggest that expectations of relevance do play a role in comprehension, but in a strictly limited way. It might be claimed, for example, that while utterances in general create expectations of truthfulness, approximations alone create expectations of relevance, which have a role to play in the case of loose talk, but only there. This account (apart from being unparsimonious) raises the following problem. As noted above, while we are all capable of realizing on reflection that an utterance was an approximation rather than a strictly literal truth, the fact that an approximation has been used is simply not noticed in the normal flow of discourse, and is surely not recognizable in advance of the comprehension process. But in that case, how could loose talk and literal talk be approached and processed with different expectations?

Here is the answer. It is not just approximations but all utterances—literal, loose or figurative—that are approached with expectations of relevance rather than truthfulness. Sometimes, the only way of satisfying these expectations is to understand the utterance as literally true. But just as an utterance can be understood as an approximation without being recognized and categorized as such, so it can be literally understood without being recognized and categorized as such. We will argue that the same is true of tropes. Literal, loose, and figurative interpretations are arrived at in the same way, by constructing an interpretation which satisfies the hearer's expectations of relevance (for earlier arguments along these lines, see Sperber and Wilson 1986a,b; 1990; 1998a).

No special machinery is needed to explain the interpretation of loose talk. In particular, contextually-determined standards of precision play no role in the interpretation process. They do not help with cases such as 'run', or 'Kleenex', which are neither semantically nor pragmatically vague; and to appeal to them in analysing cases such as 'flat' or 'at five o'clock', which might be seen as involving a pragmatic form of vagueness, would be superfluous at best.

4. Relevance: theory

Grice's maxim of truthfulness was part of what might be called an inferential model of human communication. This contrasts with a more classical code model, which treats utterances as signals encoding the messages that speakers intend to convey. On the classical view, comprehension is achieved by decoding signals to obtain the associated messages. On the inferential view, utterances are not signals but pieces of evidence about the speaker's meaning, and comprehension is achieved by inferring this meaning from the evidence provided. An utterance is, of course, a linguistically coded piece of evidence, so that the comprehension process will involve an element of decoding. But the linguistically-encoded sentence meaning need not be identical to the speaker's meaning—and we would argue that it never is—since it is likely to be ambiguous and incomplete in ways the speaker's meaning is not. On this approach, the linguistic meaning recovered by decoding is just one of the inputs to an inferential process which yields an interpretation of the speaker's meaning.

Grice, Lewis and others who have contributed to the development of an inferential approach to communication have tended to minimise the gap between sentence meaning and speaker's meaning; they treat the inference from sentence meaning to speaker's meaning as merely a matter of assigning referents to referring expressions, and perhaps of calculating implicatures. While the slack between sentence meaning and speaker's explicit meaning cannot be entirely eliminated, a framework with a maxim or convention of truthfulness has the effect of reducing it to a minimum. But why should this be something to be a priori expected or desired? Comprehension is a complex cognitive process. From a cognitive point of view, how much of the work is done by inference and how much by decoding depends on how efficient the inferential processes are. We have argued (Wilson and Sperber 1981; 1993; forthcoming; Sperber and Wilson 1986a; 1998a) that relevance-oriented inferential processes are efficient enough to allow for a much greater slack between sentence meaning and speaker's meaning than is generally assumed. Here, we summarize the theory briefly for purposes of the present discussion.

We characterize *relevance* as a property of inputs to cognitive processes which makes them worth processing. ('Relevance' is used in a technical sense which is not meant to capture any of the ordinary senses of the word.) These inputs may be external stimuli (e.g. a smell, the sound of an utterance), or internal representations which may undergo further processing (e.g. the recognition of a smell, a memory, the linguistic decoding of an utterance). At each point in our cognitive lives, there are many more potential inputs available than we can actually process: for example, we perceive many more distal stimuli than we can attend to, and have many more memories than we can reactivate at a single time. Efficiency in cognition is largely a matter of allocating our processing resources so as to maximise cognitive benefits. This involves processing inputs that offer the best expected cost/benefit ratio at the time.

Here we will consider only one type of cognitive benefit: improvements in knowledge (theoretical or practical). This is plausibly the most important type of cognitive benefit. There may be others: improvements in memory or imagination, for example (although it might be argued that these are benefits only because they contribute indirectly to improvements in knowledge; better memory and imagination lead to better non-demonstrative inference, and therefore to better knowledge). In any case, for our present purposes, there is another important reason for identifying cognitive benefits with improvements in knowledge.

In a situation where it is clear to both participants that the hearer's goal in listening to the speaker's utterances is not the improvement of knowledge—say, he just wants to be amused—there is no reason why the speaker should be expected to tell the truth. Thus, one way of challenging the maxim or convention of truthfulness would be to start by questioning whether humans are much interested in truth (e.g. Stich 1990). Here, we want to present a more pointed challenge to Grice's and Lewis's ideas, based on the nature of human communication rather than the goals of cognition. We will therefore grant that one of the goals of most human communication (though certainly not the only one) is the transmission of genuine information and the improvement of the hearer's knowledge. We will consider only cases where hearers are interested in truth. Our claim is that even in these cases, hearers do not expect utterances to be literally true.

The processing of an input in the context of existing assumptions may improve the individual's knowledge not only by adding a new piece of information, but by revising his existing assumptions, or yielding conclusions not derivable from the new piece of knowledge alone or from existing assumptions alone. We define an input as *relevant* when and only when it has such positive cognitive effects.[4] Relevance is also a matter of degree, and we want to characterize it not only as a classificatory notion but also as a comparative one. There are inputs with at least some low degree of relevance all around us, but we cannot attend to them all. What makes an input worth attending to is not just that it is relevant, but that it is more relevant than any alternative potential input to the same processing resources at that time. Although relevance cannot be measured in absolute terms, the relevance of various inputs can be compared.

For our purpose, which is to characterize a property crucial to cognitive economy, the relevance of inputs must be comparable not only in terms of benefits (i.e. positive cognitive effects), but also in terms of costs (i.e. processing effort). We therefore propose the following comparative notion of relevance:

(17) *Relevance of an input to an individual at a time*

(a) Everything else being equal, the greater the positive cognitive effects achieved in an individual by processing an input at a given time, the greater the relevance of the input to that individual at that time.

(b) Everything else being equal, the smaller the processing effort expended by the individual in achieving those effects, the greater the relevance of the input to that individual at that time.

Here is a brief and artificial illustration. Peter wakes up feeling unwell and goes to the doctor. On the basis of her examination, the doctor might make any of the following true statements:

(18) You are ill.

(19) You have flu.

(20) You have flu or 29 is the square root of 843.

The literal content of all three utterances would be relevant to Peter. However, (19) would be more relevant than either (18) or (20). It would be more relevant than (18) for reasons of cognitive effect, since it yields

[4] A positive cognitive effect is a genuine improvement in knowledge. When false information is mistakenly accepted as true, this is a cognitive effect, but not a positive one: it does not contribute to relevance (though it may seem to the individual to do so). For discussion, see the Postface to the second edition of *Relevance* (Sperber and Wilson 1995), section 3.2.1.

all the consequences derivable from (18) and more besides. This is an application of clause (a) of the characterization of relevance in (17). It would be more relevant than (20) for reasons of processing effort: although (19) and (20) yield exactly the same consequences, these consequences are easier to derive from (19) than from (20), which requires an additional effort of parsing and inference (in order to realize that the second disjunct is false and the first is therefore true). This is an application of clause (b) of the characterization of relevance in (17).

Given this characterization of relevance, it is, *ceteris paribus*, in the individual's interest to process the most relevant inputs available. We claim that this is what people tend to do (with many failures, of course). They tend to do it not because they realize that it is in their interest (and they certainly do not realize it in those terms), but because they are cognitively-endowed evolved organisms. In biological evolution, there has been constant pressure on the human cognitive system to organise itself so as to select inputs on the basis of their expected relevance (see Sperber and Wilson 2002). Hence:

(21) *The First, or Cognitive, Principle of Relevance*

The human cognitive system tends towards processing the most relevant inputs available.

The tendency described in the Cognitive Principle of Relevance is strong enough, and manifest enough, to make our mental processes at least partially predictable to others. We are in general fairly good at predicting which of the external stimuli currently affecting some other individual's nervous system she is likely to be attending to, and which of the indefinitely many conclusions that she might draw from it she will in fact draw. What we do, essentially, is assume that she will pay attention to the potentially most relevant stimulus, and process it so as to maximise its relevance: that is, in a context of easily accessible background assumptions, where the information it provides will carry relatively rich cognitive effects.

This mutual predictability is exploited in communication. As communicators, we provide stimuli which are likely to strike our intended audience as relevant enough to be worth processing, and to be interpreted in the intended way. A communicator produces a stimulus—say an utterance—which attracts her audience's attention, and she does so in an overtly intentional way. In other words, she makes it manifest that she wants her audience's attention. Since it is also manifest that the audience will tend to pay appropriate attention only to an utterance that seems relevant enough, it is manifest that the communicator wants her audience to assume that the utterance is indeed relevant enough. There is thus a minimal level of relevance that the audience is encouraged to expect: the utterance should be relevant enough to be worth the effort needed for comprehension.

Is the audience entitled to expect more relevance than this? In certain conditions, yes. The communicator wants to be understood. An utterance is most likely to be understood when it simplifies the hearer's task by demanding as little effort from him as possible, and encourages him to pay it due attention by offering him as much effect as possible. The smaller the effort, and the greater the effect, the greater the relevance. It is therefore manifestly in the communicator's interest for the hearer to presume that the utterance is not just relevant enough to be worth his attention, but more relevant than this. How much more? Here, the communicator is manifestly limited by her own abilities (to provide appropriate information, and to present it in the most efficient way). Nor can she be expected to go against her own preferences (e.g. against the goal she wants to achieve in communicating, or the rules of etiquette she wishes to follow). Still, it may be compatible with the communicator's abilities and preferences to go beyond the minimally necessary level of relevance. We define a notion of *optimal relevance* (of an utterance, to an audience) which takes these ideas into account, and propose a second principle of relevance based on it:

(22) *Optimal relevance of an utterance*

An utterance is optimally relevant to the hearer iff:

(a) It is relevant enough to be worth the hearer's processing effort;

(b) It is the most relevant one compatible with the speaker's abilities and preferences.

(23) *The Second, or Communicative, Principle of Relevance*

Every utterance conveys a presumption of its own optimal relevance.

In interpreting an utterance, the hearer invariably has to go beyond the linguistically-encoded sentence meaning. There will be ambiguities and referential indeterminacies to resolve, and other underdeterminate aspects of explicit content that we will look at shortly. There may be implicatures to identify, illocutionary indeterminacies to resolve, metaphors and ironies to interpret. All this requires an appropriate set of contextual assumptions. The Communicative Principle of Relevance and the definition of optimal relevance suggest a practical procedure for constructing a hypothesis about the speaker's meaning. The hearer should consider interpretive hypotheses (disambiguations, reference assignments, implicatures, etc.) in order of accessibility—that is, follow a path of least effort—and stop when he arrives at an interpretation which satisfies the expectations of relevance raised by the utterance itself.[5]

What makes it reasonable for the hearer to follow a path of least effort is that the speaker is expected (within the limits of her abilities and preferences) to make her utterance as easy as possible for the hearer to understand. Since relevance varies inversely with effort, the very fact that an interpretive hypothesis is easily accessible gives it an initial degree of plausibility (an epistemic advantage specific to communicated information).

What makes it reasonable for the hearer to stop at the first interpretation which satisfies his expectations of relevance is that either this interpretation is close enough to what the speaker meant, or she has failed to communicate her meaning. A speaker who produced an utterance with two or more significantly different interpretations, each yielding the expected level of cognitive effect, would put the hearer to the gratuitous and unexpected extra effort of choosing among them, and the resulting interpretation (if any) would not satisfy clause (b) of the presumption of optimal relevance.[6] Thus, when a hearer following the path of least effort finds an interpretation which satisfies his expectations of relevance, in the absence of contrary evidence, this is the best possible interpretive hypothesis. Since comprehension is a non-demonstrative inference process, this hypothesis may well be false. Typically, this hap-

[5] A hearer's expectations of relevance may be more or less sophisticated. In an unsophisticated version, presumably the one always used by young children, what is expected is actual optimal relevance. In a more sophisticated version, used by competent adult communicators who are aware that the speaker may be mistaken about what is relevant to the hearer, or in bad faith and merely intending to appear relevant, what is expected is a speaker's meaning that it may have seemed to the speaker would seem optimally relevant to the hearer. Adult communicators may nevertheless expect actual optimal relevance by default. Here we will ignore these complications, but see Sperber (1994); Wilson (2000).

[6] In the case of deliberate equivocation, where an utterance is intentionally constructed so that two apparently satisfactory competing interpretations occur to the hearer and he is unable to choose between them, neither interpretation is directly accepted. Rather, it is the fact that the speaker has produced such an utterance that is seen as a communicative act. It receives a higher-order interpretation, which may involve endorsing both lower-order interpretations (if they are compatible), or rejecting both (if they are not). For examples and discussion, see Sperber and Wilson (1987, p. 751).

pens when the speaker expresses herself in a way that is inconsistent with the expectations she herself has raised, so that the normal inferential routines of comprehension fail. Failures in communication are common enough: what is remarkable and calls for explanation is that communication works at all.

This relevance-theoretic account not only describes a psychological process but also explains what makes this process genuinely inferential: that is, likely to yield true conclusions (in this case, intended interpretations) from true premisses (in this case, from the fact that the speaker has produced a given utterance, together with contextual information). On the descriptive level, it has testable implications, some of which have been tested and confirmed by a growing body of research in experimental pragmatics.[7] On the explanatory level, it claims that what makes this relevance-guided process genuinely inferential is the fact that it typically yields a single interpretation for a given quadruple of speaker, hearer, utterance and situation. Given that such interpretations are predictable by the speaker, the best utterance for a speaker to produce is the one that is likely to be interpreted in the intended way, and the best interpretation for a hearer to choose is the one arrived at by use of the relevance-guided procedure, which is therefore likely to have been predicted and intended by the speaker. Communication is a form of co-ordination, and runs into co-ordination problems which are partly standard, and partly specific to communication. Relevance-guided comprehension takes advantage of the communication-specific aspects of these problems, and provides a solution which is, of course, imperfect, but is nonetheless effective (for further discussion, see Sperber and Wilson 2002).

5. Relevance: illustration

An utterance has two immediate effects: it indicates that the speaker has something to communicate, and it determines an order of accessibility in which interpretive hypotheses will occur to the hearer. Here is an illustration.

[7] For instance, Jorgensen, Miller and Sperber (1984); Happé (1993); Sperber, Cara, and Girotto (1995); Politzer (1996); Gibbs and Moise (1997); Hardman (1998); Matsui (1998; 2000); Nicolle and Clark (1999); Van der Henst (1999); Noveck, Bianco, and Castry (2001); Girotto, Kemmelmeir, Van der Henst, and Sperber (2001); Van der Henst, Carles, and Sperber (forthcoming); Van der Henst, Sperber, and Politzer (2002). As noted above, Van der Henst, Carles, and Sperber (forthcoming) provides a direct experimental test of the claim that considerations of relevance outweigh considerations of truthfulness in verbal communication.

Lisa drops by her neighbours, the Joneses, who have just sat down to supper:

(24) *Alan Jones*: Do you want to join us for supper?

Lisa: No, thanks. I've eaten.

A standard semantic analysis of the second part of Lisa's utterance would assign it the following truth condition:

(25) At some point in a time span whose endpoint is the time of utterance, Lisa has eaten something.

Clearly, though, Lisa means something more specific than this. She means that she has eaten that very evening, and not just anything, but a supper or something equivalent: a few peanuts wouldn't do.[8]

Here is our explanation of how Alan understands Lisa's meaning. Her utterance activates in his mind, via automatic linguistic decoding, a conceptual structure which articulates in the grammatically specified way the concepts of Lisa, of eating, and of a time span whose endpoint is the time of utterance.[9] He does not have to reason, because it is all routine, but he might reason along the following lines: she has caused me a certain amount of processing effort (the effort required to attend to her utterance and decode it). Given the Communicative Principle of Relevance, this effort was presumably not caused in vain. So the conceptual structure activated by her utterance must be a good starting point for inferring her meaning, which should be relevant to me.

Lisa's utterance, 'I have eaten', immediately follows her refusal of Alan's invitation to supper. It would be relevant to Alan (or so she may have thought) to know the reasons for her refusal, which have implications for their relationship: Did she object to the offer? Would she accept it another time? It all depends on the reasons for her refusal. The use of the perfect 'have eaten' indicates a time span ending at the time of utterance and starting at some indefinite point in the past. Alan narrows the time span by assuming that it started recently enough for the information that Lisa has eaten during that period to yield adequate consequences: here, the relevant time span is that very evening (for discussion, see Wilson and Sperber 1998). He does the same in deciding what she ate. In the circumstances, the idea of eating is most easily

[8] This is a variant of an example introduced in Sperber and Wilson (1986a, pp. 189–90)—there it was 'I have had breakfast'—which has been much discussed (e.g. Recanati (1989); Bach (1994); Carston (1998; forthcoming); Taylor (forthcoming)).

[9] We are using 'concept' in the psychological sense, to mean (roughly) the mental representation of a property.

fleshed out as eating supper, and this, together with the narrowing of the time span, yields the expected level of cognitive effect. Alan then assumes that Lisa intended to express the proposition that she has eaten supper that evening, and to present this as her reason for refusing his invitation. Although this attribution of meaning is typically a conscious event, Alan is never aware of the process by which he arrived at it, or of a literal meaning equivalent to (25).

(26a) Lisa has said to Alan, 'I have eaten'.	*Decoding of Lisa's utterance.*
(26b) Lisa's utterance is optimally relevant to Alan.	*Expectation raised by the recognition of Lisa's utterance as a communicative act, and acceptance of the presumption of relevance it automatically conveys.*
(26c) Lisa's utterance will achieve relevance by explaining her immediately preceding refusal of Alan's invitation to supper.	*Expectation raised by (b), together with the fact that such an explanation would be most relevant to Alan at this point.*
(26d) The fact that one has already eaten supper on a given evening is a good reason for refusing an invitation to supper that evening.	*First assumption to occur to Alan which, together with other appropriate premisses, might satisfy expectation (c). Accepted as an implicit premiss of Lisa's utterance.*
(26e) Lisa has eaten supper that evening.	*First enriched interpretation of Lisa's utterance as decoded in (a) to occur to Alan which might combine with (d) to lead to the satisfaction of (c). Accepted as Lisa's explicit meaning.*
(26f) Lisa is refusing supper with us because she has already had supper this evening.	*Inferred from (d) and (e), satisfying (c) and accepted as an implicit conclusion of Lisa's utterance.*
(26g) Lisa might accept an invitation to supper another time.	*From (f) plus background knowledge. One of several possible weak implicatures of Lisa's utterance which, together with (f), satisfy expectation (b).*

Table 1

The process by which Alan interprets Lisa's meaning may be represented as in Table 1, with Alan's interpretive hypotheses on the left, and his basis for arriving at them on the right. We have presented the hypotheses in English, but for Alan they would be in whatever is the medium of conceptual thought, and they need not correspond very closely to our paraphrases

We do not see this as a sequential process, starting with (26a) and ending with (26g). For one thing, interpretation is carried out 'on line,' and begins while the utterance is still in progress. Some tentative or incomplete interpretive hypotheses may be made and later revised or completed in the light of their apparent consequences for the overall interpretation. We assume, then, that interpretive hypotheses about explicit content and implicatures are developed in parallel, and stabilize when they are mutually adjusted, and jointly adjusted with the hearer's expectations of relevance.

In the present case, Alan assumes in (26b) that Lisa's utterance, decoded as in (26a), is optimally relevant to him. Since what he wants to know at this point is why she refused his invitation, he assumes in (26c) that her utterance will achieve relevance by answering this question. In this context, Lisa's utterance, 'I have eaten', provides easy access to the piece of common background knowledge in (26d)—that people don't normally want to eat supper twice in one evening. This could be used as an implicit premiss in deriving the expected explanation of Lisa's refusal, as long as her utterance is interpreted on the explicit side as conveying the information in (26e): that she has eaten supper that evening. By combining the implicit premiss in (26d) and the explicit premiss in (26e), Alan arrives at the implicit conclusion in (26f), from which further weaker implicatures, including (26g) and others, may be derived (on the notion of a weak implicature, see below). This overall interpretation satisfies Alan's expectations of relevance. On this account, explicit content and implicatures (implicit premisses and conclusions) are arrived at by a process of mutual adjustment, with hypotheses about both being considered in order of accessibility.

There is a certain arbitrariness about the way we have presented Alan's interpretive hypotheses. This is partly because, as noted above, we had to put into English thoughts which may not have been articulated in English. Another reason is that Lisa's utterance licenses not a single interpretation but any one of a range of interpretations with very similar import. By constructing any particular interpretation from this range, Alan achieves comprehension enough and has no reason to look for a better interpretation. Thus, he might take Lisa to mean either that she has had supper that evening or, more cautiously, that, whether or not what she has eaten can properly be described as supper, she has eaten enough not to want supper now. He may take her to be implicating (26g), or some conclusion similar to (26g), or nothing of the sort. In each case his interpretation is reasonable, in the sense that Lisa's utterance has encouraged him to construct it.

If Alan interprets Lisa as meaning that she has had supper, or as implicating something like (26g), he has to take some of the responsibility for the interpretation he has chosen. But this is something that hearers often do, and that speakers intend (or at least encourage) them to do. Often, the hearer will be unable to find an interpretation which is relevant in the expected way without taking some of the responsibility for it: that is, without going beyond what the speaker commits herself to acknowledging as exactly what she meant. This is typical in loose use and creative metaphor, both of which involve the communication of weak implicatures (implicatures which the hearer is given some encouragement but no clear mandate to construct). Nor is this sharing of responsibility a sign of imperfect communication: it may be just the degree of communication that suits both speaker and hearer.

Lisa's explicit meaning, as understood by Alan, logically implies the literal, unenriched meaning of her utterance: if she has eaten supper that evening, she has eaten *tout court*. Her utterance might therefore be classified as literal, for whatever good it might do. However, Alan does not attend to the literal meaning at any stage, and the fact that the utterance is literal plays no role in the communication process. This is even more obvious in the following alternative version of the dialogue:

(27) *Alan*: Do you want to join us for supper?

Lisa: I'd love to. I haven't eaten.

Here, if the literal meaning of Lisa's utterance, 'I haven't eaten', is the negation of (25) (i.e. the proposition that she has never eaten anything), then her utterance is patently false. However, this absurd interpretation never crosses Alan or Lisa's mind.

One way of avoiding such counterintuitive assignments of literal meaning would be to treat the perfect 'has eaten' as containing a hidden linguistic constituent denoting a contextually determinate time span. In (27), Lisa might then be seen as referring, via this hidden constituent, to the evening of utterance, and the fact that she has eaten plenty in her lifetime would not falsify her statement, even literally understood. We will argue below that this move is *ad hoc* and unnecessary, but let us accept it here for the sake of argument.

Assume, then, that the literal meaning of Lisa's utterance in (27) is that she has not eaten anything that evening. Now suppose that she has in fact eaten a couple of peanuts, so that her utterance is strictly speaking false. Although it may be false, it is not misleading. Rather, it is a case of loose use. Alan takes Lisa to be saying that she has not eaten supper that evening. He arrives at this interpretation by taking the concept of eating, which has been activated in his mind by automatic linguistic decoding, and narrowing it down to the concept of eating supper, which yields an overall interpretation that satisfies his expectations of relevance. The procedure is the same as for dialogue (24), but since the narrowed concept falls within the scope of a negation, the result is a loosening rather than a narrowing of the literal meaning.

It might be argued, of course, that Lisa's utterance contains a second hidden linguistic constituent denoting the food she has eaten. On this interpretation, the linguistically-determined truth-conditional meaning of 'I have eaten' is equivalent not to 'I have eaten something', but to 'I have eaten *x*', where the value of *x* (like the referent of the pronoun 'I' and the time of utterance) must be specified before the sentence token can be said to express a proposition.

In other situations, what the speaker means by saying that she has or hasn't eaten might also involve a specification of the place of eating, some manner of eating, and so on:

(28) 'I've often been to their parties, but I've never eaten anything' [*there*]

(29) 'I must wash my hands: I've eaten' [*using my hands, rather than, say, being spoon-fed*]

To deal with all such cases, one might postulate more and more hidden constituents, so that every sentence would come with a host of hidden constituents, ready for all kinds of ordinary or extraordinary pragmatic circumstances. In this way, the very idea of loose use could be altogether avoided. We see this as a *reductio* argument which goes all the way to challenging what we accepted earlier for the sake of argument: that the use of the perfect carries with it a hidden constituent denoting a given time span. There is no need to postulate such a hidden constituent: the same process which explains how 'eating' is narrowed down to 'eating supper' also explains how the time span indicated by the perfect is narrowed down to the evening of utterance. Moreover, the postulation of such hidden constituents is *ad hoc*: its role is to reduce to a minimum the slack between sentence meaning and speaker's meaning, a slack which is uncomfortable from certain theoretical viewpoints. However, we read the evidence as showing that the slack actually *is* considerable, and we adopt a theoretical viewpoint which might help us describe and understand this fact.[10]

6. The explicit communication of unencoded meanings

We are exploring the idea that the linguistically-encoded sentence meaning gives no more than a schematic indication of the speaker's meaning. The hearer's task is to use this indication, together with background knowledge, to construct an interpretation of the speaker's meaning, guided by expectations of relevance raised by the utterance itself. The conceptual resources brought to this task include all the concepts encoded in the hearer's language, but they go well beyond this (Sperber and Wilson 1998a). In particular, a concept may be recognized in context as a constituent of the speaker's explicit[11] meaning even though there is no expression in the sentence uttered, or indeed in the language, which has this concept as its linguistically-encoded meaning. This happens regularly in cases of loose use.

Consider Sue chatting to her friend Jim in the street. She looks at her watch and says:

(30) I can't stay. I must run to the bank.

The process by which Jim interprets Sue's utterance, 'I must run to the bank', may be represented as in Table 2.

What Jim takes to be Sue's explicit meaning may be described as in (31e):

(31e) Sue must RUN* to the bank (where RUN* is the meaning indicated by 'run', and is such that Sue's having to RUN* to the bank is relevant-as-expected in the context).

[10] There is now a considerable literature on hidden constituents (and more generally, on possible pragmatic contributions to explicit content). See, for example, Bach (1994; 1997; 2000); Groefsema (1995); Bezuidenhout (1997); Stainton (1997; 1998); Carston (2000; forthcoming); Neale (2000); Stanley (2000); Stanley and Szabo (2000); Taylor (forthcoming); Recanati (forthcoming).

[11] This will obviously involve some rethinking of the notion of explicitness itself. We do this in §7 below.

(31a) Sue has said to Jim, 'I must run to the bank'.	*Decoding of Sue's utterance.*
(31b) Sue's utterance is optimally relevant to Jim.	*Expectation raised by the recognition of Sue's utterance as a communicative act, and acceptance of the presumption of relevance it automatically conveys.*
(31c) Sue's utterance will achieve relevance by explaining why she must break off their chat.	*Expectation raised by (b), together with the fact that such an explanation would be most relevant to Jim at this point.*
(31d) Having to hurry to the bank on urgent business is a good reason for breaking off a chat.	*First assumption to occur to Jim which, together with other appropriate premisses, might satisfy expectation (c). Accepted as an implicit premiss of Sue's utterance.*
(31e) Sue must RUN* to the bank (where RUN* is the meaning indicated by 'run', and is such that Sue's having to RUN* to the bank is relevant-as-expected in the context).	*(Description of) the first enriched interpretation of Sue's utterance as decoded in (a) to occur to Jim which might combine with (d) to lead to the satisfaction of (c). Interpretation accepted as Sue's explicit meaning.*
(31f) Sue must break off their chat because she must hurry to the bank on urgent business.	*Inferred from (d) and (e), satisfying (c) and accepted as an implicit conclusion of Sue's utterance.*
(31g) Sue is afraid that if she stays chatting any longer, the bank may close before she gets there.	*From (f) plus background knowledge. One of several possible weak implicatures of Sue's utterance which, together with (f), satisfy expectation (b).*

Table 2

This is not, of course, a proper paraphrase (let alone a proper analysis) of Sue's meaning (as understood by Jim). The notions of *a meaning indicated by a word* and of *relevance-as-expected in a context* are not constituents of Sue's meaning, and Jim does not have to use them in understanding her utterance. As it stands, (31e) is not an interpretation but merely a description of Sue's meaning. It attributes to Sue's utterance the property of indicating rather than encoding her meaning, and to Sue's meaning the property of warranting the derivation of enough cognitive effects to make her utterance worth processing for Jim. However, it goes without saying that if Jim succeeds at all in understanding Sue's utterance, the outcome of the comprehension process will be not a description but an interpretation of Sue's meaning: that is, a mental representation which, if not identical to Sue's meaning, has a content similar enough for this to count as a case of successful comprehension. In particular, Jim's interpretation must contain an unglossed version of the concept RUN*, which on our account was not encoded but merely indicated by her use of the word 'run'.

Assuming that a satisfactory account can be given of the nature of these concepts, and of how hearers may grasp them (we will return to this below), an analysis along the lines in (31) shows how a word like 'run', or 'Kleenex', which is neither semantically nor pragmatically vague, and which (as we argued in §3 above) cannot be satisfactorily analysed by appeal to contextually-determined standards of precision, may be loosely used and understood. As we will show, the analysis is straightforwardly generalizable to the full range of cases, including 'flat' and 'five o'clock', making the appeal to contextually-determined standards of precision unnecessary.

Consider Peter and Mary discussing their next cycling trip. Peter has just said that he feels rather unfit. Mary replies:

(32) We could go to Holland. Holland is flat.

The process by which Peter interprets Mary's utterance, 'Holland is flat', may be schematically represented as in Table 3

What Peter takes to be Mary's explicit meaning may be described as in (33e):

(33e) Holland is FLAT* (where FLAT* is the meaning indicated by 'flat', and is such that Holland's being FLAT* is relevant-as-expected in the context).

As noted above, this is not an interpretation but merely a description of Mary's meaning. It attributes to Mary's utterance the property of indicating rather than encoding her meaning, and to Mary's meaning the property of warranting the derivation of enough cognitive effects to make her utterance worth processing for Peter. However, the outcome of the comprehension process must be an interpretation rather than a description of Mary's meaning. In particular, it must contain an unglossed version of the concept FLAT*, which on our account was not encoded but merely indicated by her use of the word 'flat'.

(33a) Mary has said to Peter, 'Holland is flat'.	*Decoding of Mary's utterance.*
(33b) Mary's utterance is optimally relevant to Peter.	*Expectation raised by the recognition of Mary's utterance as a communicative act, and acceptance of the presumption of relevance it automatically conveys.*
(33c) Mary's utterance will achieve relevance by giving reasons for her proposal to go cycling in Holland, which take account of Peter's immediately preceding complaint that he feels rather unfit.	*Expectation raised by (b), together with the fact that such reasons would be most relevant to Peter at this point.*
(33d) Cycling on relatively flatter terrain which involves little or no climbing is less strenuous, and would be enjoyable in the circumstances.	*First assumption to occur to Peter which, together with other appropriate premisses, might satisfy expectation (c). Accepted as an implicit premiss of Mary's utterance.*
(33e) Holland is FLAT* (where FLAT* is the meaning indicated by 'flat', and is such that Holland's being FLAT* is relevant-as-expected in the context).	*(Description of) the first enriched interpretation of Mary's utterance as decoded in (a) to occur to Peter which might combine with (d) to lead to the satisfaction of (c). Interpretation accepted as Mary's explicit meaning.*
(33f) Cycling in Holland would involve little or no climbing.	*Inferred from (d) and (e). Accepted as an implicit conclusion of Mary's utterance.*
(33g) Cycling in Holland would be less strenuous, and would be enjoyable in the circumstances.	*Inferred from (d) and (f), satisfying (b) and (c) and accepted as an implicit conclusion of Mary's utterance.*

Table 3

What might this concept FLAT* be? It is not too difficult to give a rough answer. As Mary means it, a terrain is FLAT* if travelling across it involves little or no climbing. Being FLAT* is quite compatible with small-scale unevenness, and indeed with being not plane but convex because of the curvature of the Earth.[12] However, the concept FLAT* indicated by Mary's utterance is more specific than this. It has to do with cycling when rather unfit, which determines what will count as cases of climbing. On another occasion, when travelling by car and hoping to see mountain scenery, Mary might describe the south of England as 'flat': what she would then mean is not FLAT* but some other concept FLAT**, which would be appropriately indicated in this different context by her use of the word 'flat'.

How does Peter grasp the concept FLAT* indicated by Mary's utterance? We claim that, in appropriate circumstances, the relevance-theoretic comprehension procedure should automatically guide the hearer to an acceptably close version of the concept conveyed. As noted above, the hearer's expectations of relevance warrant the assumption that the speaker's explicit meaning will contextually imply a range of specific consequences (made easily accessible, though not yet implied, by the linguistically-encoded sentence meaning). Having identified these consequences, he may then, by a process of backwards inference, enrich his interpretation of the speaker's explicit meaning to a point where it does carry these implications.

The claim that Holland is FLAT* carries a range of implications which Mary expects to satisfy Peter's expectations of relevance. The concept FLAT* is individuated (though not, of course, defined) by the fact that, in the situation described, it is the first concept to occur to Peter which determines these implications. If Mary has correctly predicted which implications Peter will actually derive from her utterance, he should arrive by a process of spontaneous backwards inference at an appropriate understanding of her explicit meaning, and in particular of the concept FLAT*.

The implications which Mary expects Peter to derive need not be individually represented and jointly listed in her mind. In normal circumstances, they would not be. She might merely expect him to derive some implications which provide reasons for going cycling in Holland, and are similar in tenor to those she herself has in mind (again without necessarily having a distinct awareness of each and every one of them). To the extent that her expectations about the implications Peter will derive are indeterminate, the same will go for the concept she intends him to arrive at by backwards inference from these implications. Notice that a difference in implications need not lead to a difference in con-

[12] In fact, for a large country, being FLAT* is incompatible with being flat: if Holland were flat, travelling from the centre to the borders would involve going upwards, that is, further away from the centre of the Earth.

cepts: from a somewhat different set of implications than the one envisaged by Mary, Peter may in fact arrive at the same concept FLAT* that she had in mind (i.e. a mental representation which picks out the same property). Suppose, however, that Peter constructs some concept FLAT** which differs slightly from FLAT*, but has roughly the same import in the situation. This would not be a case of imperfect communication or insufficient understanding. As noted above, it is quite normal for communicators to aim at such a relatively loose fit between speaker's meaning and hearer's interpretation.

We have described Mary's remark that Holland is flat as a case of loose use. We could also have described it as a case of hyperbole (i.e. as a trope). After all, taken literally, it would be a gross exaggeration. Nothing of substance hinges on whether Mary's utterance is categorized in one way or the other. Literal, loose, hyperbolical or metaphorical interpretations are arrived at by exactly same process, and there is a continuum of cases which crosscut these categories.

Consider again the case of Peter and Mary discussing their next cycling trip. Peter has just said that he feels rather unfit. In this version, Mary replies:

(34) We could go to Holland. Holland is a picnic.

This is clearly a metaphorical use of 'picnic'. The process by which Peter interprets Mary's utterance may be represented as in Table 4.

Mary uses the word 'picnic' to indicate the concept PICNIC*, which is part of what she wants to convey. Peter reconstructs this concept by treating the word 'picnic' and its associated mental encyclopaedic entry as a source of potential implicit premisses such as (35d) and (35e). From these implicit premisses and a still-incomplete interpretation of Mary's explicit meaning, he tentatively derives the implicit conclusions in (35g) and (35h), which make the utterance relevant as expected in the situation. He then arrives by backwards inference at the full interpretation of the explicit content in (35f), and its constituent concept PICNIC*.

There is an unavoidable arbitrariness about the way we have listed the implicit premisses and conclusions in (35). The more metaphorical the interpretation, the greater the responsibility the hearer has to take for the construction of implicatures (i.e. implicit premisses and conclusions), and the weaker most of these implicatures will be. Typically, poetic metaphors have a wide range of potential implicatures, and the audience is encouraged to be creative in exploring this range (a fact well recognized in literary theory since the Romantics). Communication need not fail if the implicatures constructed by the hearer are not identical to those envisaged by the speaker. Some freedom of interpretation is allowed for, and indeed encouraged, by those who speak metaphorically.

(35a) Mary has said to Peter, 'Holland is a picnic'.	*Decoding of Mary's utterance.*
(35b) Mary's utterance is optimally relevant to Peter.	*Expectation raised by the recognition of Mary's utterance as a communicative act, and acceptance of the presumption of relevance it automatically conveys.*
(35c) Mary's utterance will achieve relevance by giving reasons for her proposal to go cycling in Holland, which take into account Peter's immediately preceding complaint that he feels rather unfit.	*Expectation raised by (b), together with the fact that such reasons would be most relevant to Peter at this point.*
(35d) Going on a picnic takes little effort.	*First assumptions to occur to Peter which, together with other appropriate premisses, might satisfy expectation (c). Accepted as implicit premisses of Mary's utterance.*
(35e) Going on a picnic is a pleasant and relaxed affair.	
(35f) Holland is a PICNIC* (where PICNIC* is the meaning indicated by 'picnic', and is such that Holland's being a PICNIC* is relevant-as-expected in the context).	*(Description of) the first enriched interpretation of Mary's utterance as decoded in (a) to occur to Peter which, together with (d) and (e), might lead to the satisfaction of (c). Interpretation accepted as Mary's explicit meaning.*
(35g) Going to Holland would take little effort.	*Inferred from (d) and (f), contributing to the satisfaction of (b) and (c), and accepted as an implicit conclusion of Mary's utterance.*

Table 4

The concepts FLAT* and PICNIC* conveyed by Mary's utterances in (32) and (34) are neither encoded nor encodable in English as spoken by Mary and Peter at the time of their exchange. There is no single word or phrase of English which has FLAT* or PICNIC* as one of its linguistically-encoded senses. However, once Mary and Peter have successfully communicated one of these concepts, they may be able to co-ordinate more or less tacitly and adopt a new word or phrase to encode it, or add to the polysemy of an existing word (e.g. by giving the word 'flat' the additional stable sense FLAT*).

Different degrees of difficulty are involved in entertaining a linguistically unencodable concept such as FLAT* or PICNIC*, communicating it, and lexicalizing it. Entertaining a currently unencodable concept (i.e. a concept not encodable given the resources of the language at that time) is a relatively easy, everyday affair. As individuals, we engage in such a cognitive practice every time we discriminate and think about a property not describable by a word or phrase in our public language, which may well be several times a day. Communicating such an unencodable concept is a matter of co-ordinating the cognitive activities of two individuals so that they simultaneously attend to the same property or object. This is harder than doing it separately, but is still a relatively frequent affair. Stabilizing a word in the public language to encode such a concept involves co-ordinating cognitive dispositions in a community over time. This is much harder, and does not normally happen more than, say, a few times a year in a homogeneous speech community (see Sperber and Wilson 1998a).

7. Rethinking 'explicit,' 'literal,' and 'what is said'

If the above analysis is correct, the notions *explicit*, *literal* and *what is said*, which Grice and Lewis saw as relatively unproblematic, will have to be rethought. In this final section, we will suggest some lines on which such a rethinking might be approached.

For Grice, a speaker's meaning consists of *what is said* and (optionally) *what is implicated.* He introduced the terms 'implicate' and 'implicature' to refer to what is implicitly communicated, but rather than use the symmetrical 'explicate' and 'explicature', or just talk of what is explicitly communicated, he chose to contrast what is implicated with the ordinary-language notion *what is said*. This terminological choice reflected both a presupposition and a goal. The presupposition was that *what is said* is an intuitively clear, common-sense notion.[13] The goal was to argue against a view of meaning that ordinary-language philosophers were defending at the time. To achieve this goal, Grice wanted to

[13] See Grice (1989a, pp. 359–368) on the centrality of the intuitive notion of saying, which he characterized in what he acknowledged was 'a certain favored, and maybe in some degree artificial, sense' (Grice 1967, p. 118).

show that *what is said* is best described by a relatively parsimonious semantics, while much of the complexity and subtlety of utterance interpretation should be seen as falling on the implicit side (Carston 1998; forthcoming). We share Grice's desire to relieve the semantics of natural language of whatever can be best explained at the pragmatic level, but we take a rather different view of how this pragmatic explanation should go.

In our account, we give a theoretical status to the notions of *explicature* and *implicature* (roughly, the explicit and implicit contents of utterances), but not to the notions of *literal meaning* or *what is said*. Indeed, we introduced the term 'explicature', on the model of Grice's 'implicature', because we doubt that there is any common-sense notion of *what is said* capable of playing a useful role in the study of verbal comprehension. In our framework, explicatures are arrived at by a combination of decoding and inference, while implicatures are wholly inferred. Identifying the explicature[14] of an utterance is a matter of disambiguating, enriching and fine-tuning the semantic schema obtained by linguistic decoding. Inferring the implicatures is a matter of identifying implicit premisses and conclusions which yield an overall interpretation that is relevant in the expected way. As we have shown above, explicatures and implicatures are typically constructed in parallel, via mutual adjustment of interpretive hypotheses guided by considerations of relevance.

We have already argued that implicatures may vary in strength. The same is true of explicatures. The identification of an explicature involves a certain amount of inference. Since the inference process is non-demonstrative and draws on background knowledge, the hearer must take a certain degree of responsibility for how it comes out. How much responsibility he has to take varies from utterance to utterance: explicatures may be weaker or stronger, depending on the degree of indeterminacy introduced by the inferential aspect of comprehension. To illustrate, let us return to dialogue (24) and consider three new versions of Lisa's answer in (36a–c):

(24) (a) *Alan Jones*: Do you want to join us for supper?

(b) *Lisa:* No, thanks. I've eaten.

(36) (a) *Lisa*: No, thanks. I've already eaten supper.

[14] We are considering here only what we call basic or first-level explicatures. We also claim that there are higher-level explicatures which do not normally contribute to the truth conditions of the utterance (Wilson and Sperber 1993; Carston 2000; Ifantidou 2001).

(b) *Lisa*: No, thanks. I've already eaten tonight.

(c) *Lisa*: No, thanks. I've already eaten supper tonight.

In identifying the explicit content of all four answers, a certain amount of inference (and hence a certain degree of indeterminacy) is involved. It might be thought that the only inferences involved in (36c) are automatic—just a matter of fixing the referents of 'I' and 'tonight'—but this would be a mistake. When Lisa describes what she has eaten as supper, she may be speaking loosely. She may have had a sandwich, and be unwilling to eat again for that reason. So Alan might reasonably take 'supper' to mean SUPPER*: that is, say, enough food to be a substitute for supper. If, instead, he takes 'supper' to mean SUPPER (i.e. a regular evening meal), this is no less inferential. Whichever of the two interpretations is the first to come to mind will yield an overall interpretation which is relevant as expected, and will therefore be accepted.

Note that the first meaning to occur to Alan need not be the encoded meaning SUPPER. Suppose he knows that Lisa generally has a salad or a sandwich instead of supper: then by saying that she has eaten 'supper', she may make SUPPER* more easily accessible than SUPPER. More generally, the most accessible sense need not be the linguistically-encoded one, so when an encoded lexical sense is in fact chosen, the same process is involved as when a word is taken to convey a non-encoded sense. In each case, the first sense accessed and found to contribute to a relevant-as-expected interpretation is taken to be the intended one.

All four answers (24b) and (36a–c) communicate not just the same overall content but also the same explicature and implicatures. If this is not immediately obvious, there is a standard test for deciding whether some part of the communicated content is explicitly or implicitly conveyed. The test involves checking whether the item falls within the scope of logical operators when embedded into a negative or conditional sentence: explicatures fall within the scope of negation and other logical operators, while implicatures do not (Carston 1988; Recanati 1989; Wilson and Sperber 1998; Ifantidou 2001). Thus, consider the hypothesis that the explicature of (24b) is simply the trivial truth that Lisa has eaten at some point before the time of utterance, and that she is merely *implicating* that she has eaten that evening. The standard embedding test suggests that this hypothesis is false. If Lisa had produced the utterance, 'I haven't eaten' (as in dialogue (27)), she would clearly not have been saying that she has never eaten in her life, but merely denying that she has eaten 'supper' (i.e. SUPPER or SUPPER*) that very evening. So in saying, 'I've eaten,' Lisa is explicitly communicating that she has eaten 'supper' that very evening.

Although all four answers convey the same explicature—that Lisa has eaten SUPPER (or SUPPER*) that evening—there is a clear sense in which it is weaker in (24b) than in (36a) or (36b), and stronger in (36c). The greater the inferential element involved (and hence the greater the indeterminacy), the weaker the explicature will be. In particular, *ceteris paribus*, the greater the gap between the encoded meaning of the word and the concept conveyed by use of that word, the weaker the explicature will be. With metaphors, explicatures are at their weakest.

When the explicature is quite strong, and in particular when each word in an utterance is used to convey (one of) its encoded meaning(s), what we are calling the explicature is close to what might be commonsensically described as the explicit content, or what is said, or the literal meaning of the utterance. Whether the explicature is strong or weak, the notion of explicature applies straightforwardly. However, things go differently with the common-sense notions of *literal meaning* and *what is said*.

The notion of literal meaning, which plays such a central role in most theories of language use, is unclear in many respects. Suppose we define the literal meaning of a sentence as one of its linguistically-encoded senses. Then the literal meaning of a sentence never coincides with what the speaker explicitly communicates by uttering this sentence (except in the case of genuine 'eternal sentences', if such things exist or are ever used). A speaker's meaning is typically propositional, and at the very least, reference resolution is needed in order to get from a sentence meaning to a proposition. It seems more appropriate, then, to define the literal meaning of an utterance (rather than a sentence) as the proposition obtained by combining its linguistic sense with its reference. When the speaker's meaning coincides with this proposition, we do indeed have a prototypical case of literalness. Suppose an anthropologist confesses:

(37) I have eaten human flesh.

In most situations, (37) would be relevant enough if it were understood as explicitly communicating its literal meaning, without any narrowing of the time span or the way in which the eating of human flesh is understood to have taken place. This is then a prototypical case where literal meaning (understood as sense plus reference) coincides with the speaker's explicit meaning, or explicature. However, such cases are the exception rather than the rule.

In the first place, there are cases where the explicature cannot simply correspond to the combination of linguistic sense plus reference, because this is not enough to determine a unique proposition. Consider (38):

(38) His car is too big.

Even when the linguistic sense is combined with appropriate referents for the pronoun 'his' and the present tense, the result is not a complete proposition.[15] 'His car' might be the car he owns, the car he is renting, the car he is thinking about, and so on, and deciding which it is meant to be is not a matter of disambiguation or reference assignment, but of enriching the linguistically-encoded meaning. Similarly, 'too big' is indeterminate unless some contextual criterion is supplied for deciding what counts as big enough in this case. Such cases are sometimes dealt with by redefining the literal meaning of an utterance as determined by a combination of sense, plus reference, plus obligatory enrichment (sometimes known as the 'minimal proposition' expressed by an utterance).[16] Suppose, then, that (38) is enriched as in (39):

(39) The car Bob is planning to steal is too big to hide in the lorry.

Is this the literal meaning of (38) on that occasion, or is there some other, simpler literal meaning? If so, what is it? In such cases, intuitions about literalness become quite unclear.

Even leaving aside the problem of obligatory enrichment (and other related problems discussed in Searle 1979, Ch.5), and considering only sentences where the combination of sense plus reference determines a complete proposition, the fact is that in most cases, the explicature of an utterance goes well beyond this. The identification of an explicature may involve enrichment of the encoded meaning, loosening of the encoded meaning, or some combination of enrichment and loosening. Such cases are sometimes dealt with by drawing a distinction between *literal meaning* and *literal use*, and treating an utterance as a case of lit-

[15] Whatever the proposition expressed by a literal utterance of (38), it entails the existentially quantified proposition *There is a relationship between the referent of 'his' and a unique car, and there is a criterion of size, such that this car is too big by this criterion.* However, this proposition is never the utterance meaning of (38), and it would be highly counter-intuitive to treat it as its literal meaning.

[16] See, for example, Travis (1985); Carston (1988; 1998; 2000; forthcoming); Recanati (1989; forthcoming); Bach (1994; 1997; 2000).

eral use provided that the explicature departs from the literal meaning only by being richer or more specific (e.g. Katz 1990, pp 144–6). So when Lisa says, 'I have eaten', in response to Alan's invitation to supper, the literal meaning of her utterance (the proposition that Lisa has eaten at some point in a time span ending at the time of utterance) is determined in the regular way by combining sense with reference. Since her actual meaning (that she has eaten supper that evening) is an enrichment of the literal meaning, this would count as a case of literal use.[17]

This proposal also runs into problems. If enrichment of meaning preserves literalness of use, then (40) must be treated as a case of literal use:

(40) [*Antony praising Brutus in* Julius Caesar] This was a man!

However, in classical rhetoric, (40) would be classified as a case of figurative use (more precisely, as a variety of synecdoche). Here again, intuitions are probably not clear enough to decide, so the decision would have to be made on theoretical grounds. But we have argued that a notion of literalness has no role to play in a theory of language use. The interpretation of every utterance involves a process of meaning construction, which is the same whether the result is an enriched, loosened, enriched-and-loosened, or literal interpretation. Yes, literalness can be defined, or at least characterized, in terms of a prototype, but, no, verbal understanding does not involve paying any attention to literal meaning, let alone to literal use. There is no theoretical basis for sharpening our characterization of literalness. On the other hand, as we will see, there may be social pressure to do so.

Similar problems arise with the notion of *what is said*. Given that a speaker has produced some utterance *U* as an act of verbal communication, what is the proper completion of (41)?

(41) The speaker said that …

The idea that there is a theoretically adequate and useful notion of *what is said* implies that there is a correct completion of (41) (or a set of semantically equivalent completions) which uniquely captures what is said by uttering *U*. (This is, of course, compatible with recognizing that different completions may be pragmatically acceptable in different situ-

[17] Note that, in order to ground a clear notion of literalness, the notion of enrichment itself must be properly defined. It cannot be defined in terms of entailment since, presumably, literalness of use is maintained under negation or embedding (e.g. in the antecedent of a conditional), whereas entailment relations are not: 'I have eaten' is entailed by 'I have eaten supper tonight', but 'I have not eaten' is not entailed by 'I have not eaten supper tonight'; similarly with 'If I have/haven't eaten (supper tonight), then P)'.

ations. For instance, the dots might be replaced by an exegesis, a summary or a sarcastic rendering of *U*. However, these would not fit the intended notion of *what is said*.) Prototypical instances of the intended notion are easy to find: they are the same as the prototypical instances of *literal meaning*. Thus, what was said by the speaker of (37) is unproblematically rendered as (42):

(42) The anthropologist said that she has eaten human flesh.

Here, the speaker's explicit meaning can be straightforwardly rendered by a transposition from direct to indirect quotation. However, this is not always so.

When Lisa produces the utterance in (24b), 'I have eaten' (with an explicature rather weaker than the one conveyed by (36c), 'I have already eaten supper tonight'), what is she saying? Intuitions typically waver. *Saying* is often understood in an indirect-quotational sense, where what is said is properly rendered by an indirect quotation of the original utterance. It might thus be claimed that in uttering (24b), Lisa is merely saying that she has eaten; but this would not adequately capture her meaning as we have described it above. *Saying* can also be understood in a commitment sense, where what is said is what the speaker is committing herself to in producing an utterance. This is typically the sense invoked when, precisely, the competence or sincerity of the speaker's commitment is being challenged. Suppose Alan replies to Lisa's utterance (24b), 'What you just said is false: I happen to know that you haven't eaten a thing since lunch.' By common-sense standards, Alan is not misusing the word 'said'. However, his response makes sense only if he was taking Lisa to have said not just that she has eaten, but that she has eaten that very evening. Of course, Lisa might then reply that she had so much lunch that she didn't feel like eating anything more that day. While this might be seen as disingenuous, the explicature of her utterance is weak enough to leave room for reasonable doubt. On the other hand, if she had not eaten for days, then in uttering (24b) in this situation, she would undoubtedly be saying something false.

The weaker the explicature of an utterance, the harder it is to paraphrase what the speaker was saying except in the indirect-quotational sense. It is always safe to quote the speaker's words (either directly or indirectly), but this is of limited use. It would be more useful to paraphrase the speaker's meaning, except for the element of arbitrariness involved. This vacillation between a quotational and a commitment sense of 'saying' is particularly obvious in the case of metaphor. On the one hand, we may feel that here the only safe sense is the quotational one. When Mary produces the utterance in (34), 'Holland is a picnic', we would all immediately agree that she is saying that Holland is a picnic. However, this does not provide a truth-evaluable content which can be crisply and confidently paraphrased and which Mary is clearly committing herself to. On the other hand, it is quite possible to disagree with what is being said (in an everyday sense) without being able to paraphrase it. If Peter disagrees with Mary's utterance, 'Holland is a picnic', he may well tell her, 'What you say is false.' Here, he would be expressing disagreement with Mary's explicature, however vague, rather than making the obvious point that Holland is a country and not a social event.

Speakers commit themselves, and they can be criticized for their commitments. Often, however, the exact character of their commitment can be disputed. This happens quite regularly at home, in public life, and in court. Arguing about what was said—both its content and its truthfulness—is a social practice conducted within the framework of 'folk-linguistics.' The notions of *literal meaning* and *what is said* come from folk-linguistics, and they may well play a useful (or even indispensable) role within this framework. Most people are more interested in the norms governing linguistic communication than in the mechanisms by which it is achieved. The apparent platitudes listed at the beginning of this paper—as speakers, we expect what we say to be accepted as true, as hearers, we expect what is said to us to be true—are versions of one of these folk-linguistic norms, a norm of truthfulness in what is said.

In the situations where it is typically appealed to, the norm of truthfulness is a reasonable requirement on verbal communication. It is generally invoked when the audience suspects that it is being violated, and it is very rare for a speaker accused of violating it to dispute its applicability. By contrast, disagreements about what was actually said are not rare at all. The notion of literal meaning is typically invoked in the context of such disputes. It is often easier to agree on the literal meaning of an utterance, and on its literal truth or falsity, than on what the speaker meant, or what the hearer could reasonably have understood. A speaker can retreat behind the literal meaning of her utterance, which may have been true even if the utterance was misleading. A hearer can point out that what was literally said was false, and the speaker may reply that she was not intending to be taken so literally. Many such arguments are never settled. This shows the limitations of any description of the speaker's commitment in terms of the folk-linguistic notion of *saying*.

The very idea that what a speaker says should always (with the possible exception of poetry) be either literal or paraphrasable by means of a literal utterance is an illusion of folk-linguistics. Western folk-linguistics, at least, is committed to a code model of communication from which it follows that what is said should always be transparent or paraphrasable. Efforts to bring communicative practice into line with this ideal have had some effect on language use. In forms of verbal interaction where the speaker's commitments are particularly important from a social point of view (in science or law, for example), there is a demand that speech should in general be literal, and that occasional departures from literalness should be overt and obvious: occasional metaphors are acceptable, but not the loose uses found in ordinary exchanges. How well the demand is actually satisfied is another matter. In general, folk-linguistic theories about communicative practice have rather limited and peripheral effects on the natural processes of speech and comprehension, where so many of the sub-processes involved are automatic and impenetrable (cf. Levelt 1989).

It may have seemed reasonable to philosophers such as Paul Grice or David Lewis to base their philosophy of language on a reformulated norm of truthfulness. However, their reformulations did not go far enough. Both Grice and Lewis took for granted that truthfulness based on the conventional meaning of utterances is expected (for Grice, conventional meaning is just literal meaning; for Lewis, it is literal or figurative meaning, with figurative meaning being derived from literal meaning). This assumption played a central role in Lewis's explanation of how linguistic meaning could be conventional, and in Grice's account of how non-conventional meanings could be conveyed.[18]

We agree that, at least in most cases, a hearer who attends to an utterance expects to be informed of something. We agree with Grice that 'false information is not an inferior kind of information; it just is not information' (Grice 1989a, p. 371). So, yes, hearers expect to be provided with true information. But there is an infinite supply of true information which is not worth attending to. Actual expectations are of relevant information, which (because it is information) is also true.[19] However, we have argued that there just is no expectation that the true information communicated by an utterance should be literally or convention-

[18] Although her notion of a convention, and of the role of intention in communication, differs from those of Lewis or Grice, Ruth Millikan (1984) also bases her philosophy of language on a version of the norm of truthfulness (see Origgi and Sperber 2000 for discussion).

[19] For further discussion of the relation between truth and relevance, see the Postface to the second edition of *Relevance* (Sperber and Wilson 1995), section 3.2.1.

ally expressed, as opposed to being explicated or implicated in the sense we have discussed here.

Linguistically-encoded meaning is far too schematic and fragmentary to be capable of being true or false: it is just an input to further processing. Contrary to the standard view, this further processing does not consist simply in combining linguistic sense with contextual reference in order to determine a literal meaning. The fact that the speaker has produced this utterance with this linguistic meaning is expected to provide a relevant piece of evidence and a point of departure for inferring the speaker's meaning. The resulting explicatures and implicatures are in turn expected to provide worthwhile input for further processing: that is, to be relevant (and therefore true).[20]

University College London DEIRDRE WILSON
Gower Street
London WC1E 6BT, UK
deirdre@ling.ucl.ac.uk

Institut Jean Nicod (CNRS & EHESS) DAN SPERBER
1 bis Avenue de Lowendal
Paris 75007, France

dan@sperber.com

References

Bach, Kent 1994: 'Conversational Impliciture'. *Mind and Language*, 9, pp. 124–62.
——1997: 'The Semantics–Pragmatics Distinction: What it is and Why it Matters'. *Linguistische Berichte*, 8 (Special Issue on Pragmatics), pp. 33–50. Reprinted in Turner 1999.
——2000: 'Quantification, Qualification and Context: A Reply to Stanley and Szabo'. *Mind and Language*, 15, pp. 262–83.
Bezuidenhout, Anne 1997: 'Pragmatics, Semantic Underdetermination and the Referential–Attributive Distinction'. *Mind*, 106, pp. 375–409.
Carruthers, Peter and Boucher, Jill (eds) 1998: *Language and Thought: Interdisciplinary Themes*. Cambridge: Cambridge University Press.

[20] We are grateful to two anonymous reviewers for *Mind*, to Robyn Carston, Steven Gross, Corinne Iten, Pierre Jacob, Stephen Neale, Gloria Origgi, Francois Recanati, Rob Stainton, Jason Stanley, and the members of the UCL Pragmatics Reading Group for many helpful comments and suggestions.

Carruthers, Peter and Chamberlain, Andrew (eds) 2000: *Evolution and the Human Mind: Language, Modularity and Social Cognition.* Cambridge: Cambridge University Press.

Carston, Robyn 1988: 'Implicature, Explicature and Truth-Theoretic Semantics', in Kempson 1988, pp. 155–81. Reprinted in Davis 1991.

——1998: *Pragmatics and the Explicit–Implicit Distinction.* University College London Ph.D. thesis.

——2000: 'Explicature and Semantics'. *UCL Working Papers in Linguistics*, 12, 1–44. Forthcoming in Davis and Gillon.

——forthcoming: *Thoughts and Utterances: The Pragmatics of Explicit Communication.* Oxford: Blackwell.

Carston, Robyn and Uchida, Seiji (eds) 1998: *Relevance Theory: Applications and Implications.* Amsterdam: John Benjamins.

Cole, Peter (ed.) 1981: *Radical Pragmatics.* New York: Academic Press.

Cole, Peter and Morgan, Jerry 1975: *Syntax and Semantics 3: Speech Acts.* New York: Academic Press.

Davis, Steven (ed.) 1991: *Pragmatics: A Reader.* New York: Oxford University Press.

Davis, Steven and Gillon, Brendan (eds) forthcoming: *Semantics: A Reader.* New York: Oxford University Press.

Gibbs, Ray and Moise, Jessica 1997: 'Pragmatics in Understanding What is Said'. *Cognition*, 62, pp. 51–74.

Girotto, Vittorio, Kemmelmeir, Markus, Van der Henst, Jean-Baptiste, and Sperber, Dan 2001: 'Inept Reasoners or Pragmatic Virtuosos? Relevance and the Deontic Selection Task'. *Cognition*, 81, pp. B69–B76.

Grice, H. Paul 1967: *Logic and Conversation* (William James Lectures, Harvard University), in Grice 1989b, pp. 3–143. Originally published in Cole and Morgan 1975.

——1989a: 'Retrospective Epilogue', in Grice 1989b, pp. 339–85.

——1989b: *Studies in the Way of Words.* Cambridge, MA: Harvard University Press.

Groefsema, Marjolein 1995: 'Understood arguments: A semantic/pragmatic approach'. *Lingua*, 96, pp. 139–61.

Gross, Steven 1998: *Essays on Linguistic Context-Sensitivity and its Philosophical Significance.* Harvard University Ph.D. dissertation.

Gunderson, Keith (ed.) 1975: *Minnesota Studies in the Philosophy of Science*, volume VII. Minnesota: University of Minnesota Press.

Happé, Francesca 1993: 'Communicative Competence and Theory of Mind in Autism: A Test of Relevance Theory'. *Cognition*, 48, pp. 101–19.

Hardman, David 1998: 'Does Reasoning Occur in the Selection Task? A Comparison of Relevance-Based Theories'. *Thinking and Reasoning*, 4, pp. 353–76.
Horn, Laurence and Ward, Gregory (eds) forthcoming: *Handbook of Pragmatics*. Oxford: Blackwell.
Ifantidou, Elly 2001: *Evidentials and Relevance*. Amsterdam: John Benjamins.
Jorgensen, Julia, Miller, George, and Sperber, Dan 1984: 'Test of the Mention Theory of Irony'. *Journal of Experimental Psychology: General*, 113, pp. 112–20.
Katz, Jerrold 1990: *The Metaphysics of Meaning*. Cambridge, MA: MIT Press.
Keefe, Rosanna and Smith, Peter (eds) 1997: *Vagueness: A Reader*. Cambridge, MA: MIT Press.
Kempson, Ruth (ed.) 1988: *Mental Representation: The Interface between Language and Reality*. Cambridge: Cambridge University Press.
Khalfa, Jean (ed.) 1994: *What is Intelligence?* Cambridge: Cambridge University Press.
Levelt, Willem 1989: *Speaking: From Intention to Articulation*. Cambridge, MA: MIT Press.
Lewis, David 1970: 'General Semantics', in Lewis 1983, pp. 189–232. Originally published in *Synthese*, 22, 1970.
——1975: 'Languages and Language', in Lewis 1983, pp. 163–88. Originally published in Gunderson 1975.
——1979: 'Scorekeeping in a Language Game', in Lewis 1983, pp. 233–49. Originally published in *Journal of Philosophical Logic*, 8, 1979.
——1983: *Philosophical Papers*, volume I. Oxford: Oxford University Press.
Lyons, John 1977: *Semantics*, volume II. Cambridge: Cambridge University Press.
Matsui, Tomoko 1998: 'Pragmatic Criteria for Reference Assignment: A Relevance-Theoretic Account of the Acceptability of Bridging'. *Pragmatics and Cognition*, 6, pp. 47–97.
——2000: *Bridging and Relevance*. Amsterdam: John Benjamins.
Millikan, Ruth 1984: *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
Neale, Stephen 1992: 'Paul Grice and the Philosophy of Language'. *Linguistics and Philosophy*, 15, pp. 509–59.
——2000: 'On being Explicit: Comments on Stanley and Szabo, and on Bach'. *Mind and Language*, 15, pp. 284–94.

Nicolle, Steven and Clark, Billy 1999: 'Experimental Pragmatics and What is Said: A Response to Gibbs and Moise'. *Cognition*, 66, pp. 337–54.

Noveck, Ira, Bianco, Maryse, and Castry, Alain 2001: 'The Costs and Benefits of Metaphor'. *Metaphor and Symbol*, 16, pp. 109–21.

Origgi Gloria and Sperber, Dan 2000: 'Evolution, communication and the proper function of language'. In Carruthers and Chamberlain 2000, pp. 140–69.

Politzer, Guy 1996: 'A Pragmatic Account of a Presuppositional Effect'. *Journal of PsycholinguisticResearch*, 25, 543–51.

Recanati, François 1989: 'The Pragmatics of What is Said'. *Mind and Language*, 4, pp. 295–329. Reprinted in Davis 1991.

——forthcoming: 'Unarticulated constituents'. Forthcoming in *Linguistics and Philosophy.*

Searle, John 1979: *Expression and Meaning*. Cambridge: Cambridge University Press.

Sperber, Dan 1994: 'Understanding Verbal Understanding', in Khalfa 1994, pp. 179–98.

——(ed.) 2000: *Metarepresentations: A Multidisciplinary Perspective.* Oxford: Oxford University Press.

Sperber, Dan, Cara, Francisco, and Girotto, Vittorio 1995: 'Relevance Theory Explains the Selection Task'. *Cognition*, 57, pp. 31–95.

Sperber, Dan and Wilson, Deirdre 1981: 'Irony and the Use–Mention Distinction'. In Cole 1981, pp. 295–318. Reprinted in Davis 1991.

——1986a: *Relevance: Communication and Cognition.* Blackwell, Oxford, second edition 1995.

——1986b: 'Loose talk'. *Proceedings of the Aristotelian Society* LXXXVI, 1985–6, pp. 153–71. Reprinted in Davis 1991.

——1987: 'Presumptions of Relevance'. *Behavioral and Brain Sciences,* 10, pp. 736–54.

——1990: 'Rhetoric and Relevance'. In Wellbery and Bender 1990, pp. 140–55.

——1995: 'Postface' to the second edition of Sperber and Wilson 1986a, pp. 255–79.

——1998a: 'The Mapping between the Mental and the Public Lexicon', in Carruthers and Boucher 1998, pp. 184–200.

——1998b: 'Irony and Relevance: A Reply to Drs Seto, Hamamoto and Yamanashi', in Carston and Uchida 1998, pp. 283–93.

——2002: 'Pragmatics, Modularity and Mindreading'. *Mind and Language,* 17, pp. 3–23.

Stainton, Robert 1997: 'Utterance Meaning and Syntactic Ellipsis'. *Pragmatics and Cognition*, 5, pp. 51–78.
——1998: 'Quantifier Phrases, "Meaningfulness in Isolation", and Ellipsis'. *Linguistics and Philosophy*, 21, pp. 311–40.
Stanley, Jason 2000: 'Context and Logical Form'. *Linguistics and Philosophy*, 23, pp. 391–434.
Stanley, Jason and Szabo, Zoltan 2000: 'On Quantifier Domain Restriction'. *Mind and Language*, 15, pp. 219–61.
Stich, Stephen 1990: *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation.* Cambridge, MA: MIT Press.
Taylor, Ken forthcoming: 'Sex and breakfast'. Forthcoming in *Synthese.*
Travis, Charles 1985: 'On What is Strictly Speaking True'. *Canadian Journal of Philosophy*, 15, pp. 187–229.
Turner, Ken (ed.) 1999: *The Semantics–Pragmatics Interface from Different Points of View.* Amsterdam: Elsevier Science.
Unger, Peter 1975: *Ignorance: A Case for Scepticism.* Oxford: Clarendon Press.
Van der Henst, Jean-Baptiste 1999: 'The Mental Model Theory and Spatial Reasoning Re-Examined: The Role of Relevance in Premise Order'. *British Journal of Psychology*, 90, pp. 73–84.
Van der Henst, Jean-Baptiste, Carles, Laure, and Sperber, Dan forthcoming: 'Truthfulness and Relevance in Telling the Time'. Forthcoming in *Mind and Language.*
Van der Henst, Jean-Baptiste, Sperber, Dan and Politzer, Guy 2002: 'When is a Conclusion Worth Deriving? A Relevance-Based Analysis of Indeterminate Relational Problems'. *Thinking and Reasoning*, 8, pp. 1–20.
Wellbery, David and Bender, John (eds) 1990: *The Ends of Rhetoric: History, Theory, Practice.* Stanford, CA: Stanford University Press.
Werth, Paul (ed.) 1981: *Conversation and Discourse.* London: Croom Helm.
Williamson, Timothy 1994: *Vagueness.* London: Routledge.
Wilson, Deirdre and Sperber, Dan 1981: 'On Grice's Theory of Conversation', in Werth 1981, pp. 155–78.
——1992: 'On Verbal Irony'. *Lingua*, 87: 53–76.
——1993: 'Linguistic Form and Relevance'. *Lingua*, 90: 1–25.
——1998: 'Pragmatics and Time', in Carston and Uchida 1998, pp. 1–22. Amsterdam: John Benjamins.
——forthcoming: 'Relevance Theory', in Horn and Ward forthcoming.
Wilson, Deirdre 2000: 'Metarepresentation in Linguistic Communication', in Sperber 2000, pp. 411–48.