

Linguistic form and relevance

Deirdre Wilson^a and Dan Sperber^b

^a*Department of Phonetics and Linguistics, University College London, Gower Street, London WC1E 6BT, UK*

^b*C.N.R.S. and CREA, Ecole Polytechnique, 1 rue Descartes, F-75005 Paris, France*

1. Introduction

Our book *Relevance* (Sperber and Wilson 1986) treats utterance interpretation as a two-phase process: a modular decoding phase is seen as providing input to a central inferential phase in which a linguistically encoded logical form is contextually enriched and used to construct a hypothesis about the speaker's informative intention. *Relevance* was mainly concerned with the inferential phase of comprehension: we had to answer Fodor's challenge that while decoding processes are quite well understood, inferential processes are not only not understood, but perhaps not even understandable (see Fodor 1983). Here we will look more closely at the decoding phase and consider what types of information may be linguistically encoded, and how the borderline between decoding and inference can be drawn.

It might be that all linguistically encoded information is cut to a single pattern: all truth conditions, say, or all instructions for use. However, there is a robust intuition that two basic types of meaning can be found. This intuition surfaces in a variety of distinctions: between *describing* and *indicating*, *stating* and *showing*, *saying* and *conventionally implicating*, or between *truth-conditional* and *non-truth-conditional*, *conceptual* and *procedural*, or *representational* and *computational* meaning. In the literature, justifications for these distinctions have been developed in both strictly linguistic and more broadly cognitive terms.

The linguistic justification goes as follows (see for example Récanati 1987). Utterances express propositions; propositions have truth conditions; but the meaning of an utterance is not exhausted by its truth conditions, i.e. the truth conditions of the proposition expressed. An utterance not only expresses a proposition but is used to perform a variety of speech acts. It can thus be expected to encode two basic types of information: truth-conditional and

non-truth-conditional, or propositional and illocutionary – that is, information about the state of affairs it describes, and information indicating the various speech acts it is intended to perform.

The cognitive justification goes as follows (see for example Sperber and Wilson 1986, Blakemore 1987, 1992). Linguistic decoding provides input to the inferential phase of comprehension; inferential comprehension involves the construction and manipulation of conceptual representations. An utterance can thus be expected to encode two basic types of information: representational and computational, or conceptual and procedural – that is, information about the representations to be manipulated, and information about how to manipulate them.

It is tempting to assume that these two approaches are equivalent, and classify the data in identical ways. This would be so, for example, if any construction which contributed to the truth conditions of an utterance did so by encoding concepts, while all non-truth-conditional constructions encoded procedural information. We want to argue that this assumption is false. The two distinctions cross-cut each other: some truth-conditional constructions encode concepts, others encode procedures; some non-truth-conditional constructions encode procedures, others encode concepts. This raises a more general question. What is the relation between the two approaches? Are the distinctions drawn by one approach somehow more basic than those drawn by the other? For example, is it possible to predict whether a given construction is truth-conditional or non-truth-conditional on the basis of some systematic interaction between the type of information it encodes and other linguistic or cognitive factors? We will touch briefly on these issues towards the end.

These internal questions about the decoding phase of comprehension are mainly of interest to linguistic semantics. Pragmatic theorists are more interested in an external question: how is the borderline between decoding and inference to be drawn? Linguistic decoding is not the only source of input to inferential comprehension. When Peter notices Mary's accent and decides that she is Scottish, this information is not *encoded* in her utterance, any more than it is encoded by the fact that she is drinking malt whisky or wearing a Black Watch tartan kilt. These are facts about her which Peter may notice, and from which he may draw inferences. How do these inferences interact with linguistically encoded information? How do we decide, as theorists, which information was decoded and which was inferred?

In *Relevance* we tried to answer some of these questions; answers to others will be attempted here. In figure 1, we have tried to draw the threads together

and give a general picture of the various types of information, decoded and inferred, that an utterance can convey.

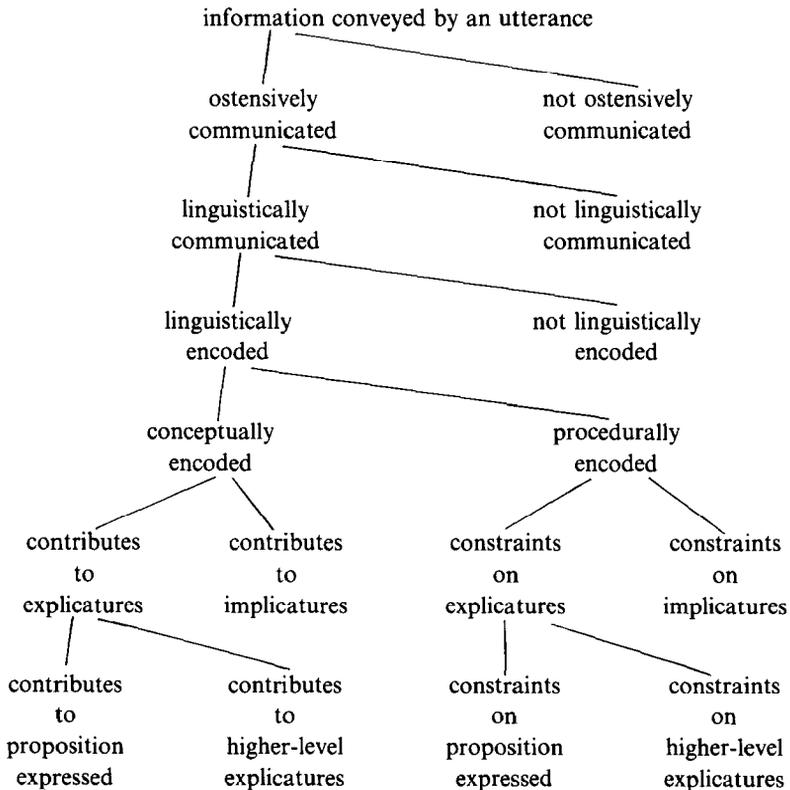


Fig. 1. Types of communicated information.

This paper is organised around the distinctions drawn in the diagram. We will start at the top, with the inferential phase of comprehension, and work down through external questions about the borderline between decoding and inference, to end with internal questions about the decoding phase.

2. Conveying and ostensively communicating

An utterance makes manifest a variety of assumptions: the hearer attends to as many of these as seem relevant to him. All these assumptions are

conveyed by the utterance. Not all of them are ostensively communicated, as the following examples will show:

- (a) Mary speaks to Peter: something in her voice or manner makes him think that she is sad. As she speaks, he is wondering about the reasons for her sadness. This is not what Mary wanted: she was trying to hide her feelings from him. In the terms of *Relevance*, Mary had neither an informative nor a communicative intention. The case is one of accidental information transmission.
- (b) Mary speaks sadly to Peter. She intends him to notice her sadness, but to think she is bravely hiding her pain. In the terms of *Relevance*, she intends to inform Peter of her sadness, but she wants her informative intention to be fulfilled without being recognised. Some form of covert (and hence non-ostensive) communication is taking place.
- (c) Mary speaks sadly to Peter. She intends him to notice her sadness, and to realise that she intended him to notice it, but to think she wanted this higher-order intention to remain hidden from him. In the terms of *Relevance*, Mary intends to inform Peter of her sadness, and she wants her informative intention to be recognised but not to become mutually manifest. Again, some form of covert communication is taking place.
- (d) Mary speaks sadly to Peter. She intends to inform him of her sadness, and she wants her informative intention to be not merely recognised, but to become mutually manifest. In the terms of *Relevance*, Mary has both an informative and a communicative intention. Ostensive communication is taking place.

In *Relevance*, we showed how examples (a)–(d) all fall within the scope of a relevance-based theory of cognition. As Mary speaks, Peter will pay attention to any aspect of her behaviour that seems relevant to him. Sometimes, to explain her behaviour, he will be led to attribute to her an informative intention. What distinguishes ostensive communication from other forms of intentional or unintentional information transmission is that the hearer has special help in recognising the speaker's informative intention. Ostensive communication creates a presumption of relevance and falls under the principle of relevance. Of all accessible hypotheses about the speaker's informative intention, the hearer should accept the first one tested and found consistent with the principle of relevance. Having recognised the speaker's informative intention by use of this criterion, he is entitled to treat it as not only manifest but mutually manifest.

3. Linguistic and non-linguistic communication

When Mary speaks sadly to Peter, intending to communicate that she is sad, his knowledge of language does not help him to recognise her informative intention. Mary communicates her sadness to Peter, but she does not *linguistically* communicate it. For an assumption to be linguistically communicated, the linguistic properties of the utterance must help with its recovery. In this example, they do not.

This is not to say that paralinguistic clues such as tone of voice or manner play no role at all in linguistic communication. Consider the exchange in (1):

- (1a) *Peter*: Can you help?
- (1b) *Mary* (sadly): I can't.

Suppose that in saying (1b), Mary expected Peter not only to notice that she is sad, but to ask himself why she is sad, and to come to the conclusion in (2):

- (2) Mary is sad that she can't help Peter.

Suppose, moreover, that Mary intended not merely to inform Peter of (2) but to communicate it ostensively. Then in the terms of *Relevance*, (2) would be an *explicature* of (1b).

An utterance typically has several explicatures. Mary's utterance in (1b) might include those in (3):

- (3a) Mary can't help Peter to find a job.
- (3b) Mary says she can't help Peter to find a job.
- (3c) Mary believes she can't help Peter to find a job.
- (3d) Mary regrets that she can't help Peter to find a job.

The explicatures of an utterance are constructed by enriching a linguistically encoded logical form to a point where it expresses a determinate proposition, such as (3a), and optionally embedding it under a higher-level description: for example, a speech-act description such as (3b), or a propositional attitude description such as (3c) or (3d). Let us call (3a) the proposition expressed by the utterance and (3b–d) its higher-level explicatures. Then not only the proposition expressed by the utterance but also all its higher-level explicatures are linguistically communicated. We will return to this point below.

Explicatures, like implicatures, have their own truth conditions, and are capable of being true or false in their own right. However, only the proposition expressed is normally seen as contributing to the truth conditions of the associated utterance. Here we will follow the standard semantic practice of calling a construction truth-conditional if and only if it contributes to the proposition expressed. This point will be important in later sections.

4. Linguistic communication and encoding

Not everything that is linguistically communicated is linguistically *encoded*. An interpretation is encoded when it is stipulated in the grammar. Since Grice's *William James Lectures* (reprinted in Grice 1989), a sustained and largely successful attack on unreflective appeals to encoding, the borderline between linguistic communication and linguistic encoding has been a major focus of pragmatic research. To illustrate recent developments in this area, we will consider some post-Gricean analyses of 'and'.

Grice showed that differences in the interpretation of conjoined utterances such as (4a) and (4b) can be explained without appeal to lexical encoding:

(4a) Peter got angry and Mary left.

(4b) Mary left and Peter got angry.

The temporal connotations of (4a) and (4b) arise not, he said, from an extra, temporal sense of 'and', but from an interaction between the regular non-temporal sense and a pragmatic maxim of orderliness which instructs speakers to recount events in the order in which they happened. In other words, the temporal connotations of (4a) and (4b) are linguistically communicated without being linguistically encoded.

There are problems with Grice's account. In the first place, (4a) and (4b) have not only temporal but causal connotations: (4a) suggests that Mary left because Peter got angry and (4b) suggests the reverse. These suggestions do not follow from a maxim of orderliness alone. Or consider (5a–d):

(5a) Peter went into the kitchen and found Mary.

(5b) Peter took out his key and opened the door.

(5c) Mary injured her leg and sued Peter.

(5d) Mary is English and cooks well.

(5a) suggests that Peter found Mary in the kitchen, (5b) that Peter used his key to open the door, (5c) that Mary sued Peter for the injury to her leg, and (5d) that she cooks well despite the fact that she is English. None of these suggestions is linguistically encoded, as witness the fact that all are cancellable without contradiction. The problem raised by such suggestions is this. Either new maxims are needed to explain them – in which case Grice's framework is incomplete. Or they are explainable in terms of existing maxims such as the maxim of relevance – in which case the temporal connotations of (4a) and (4b) should be similarly explainable, and the maxim of orderliness is redundant.¹

Relevance theory suggests the latter response. In processing (5b), for example, the hearer is looking for an interpretation consistent with the principle of relevance: typically, an interpretation which yields adequate effects for the minimum justifiable effort in a way the speaker could manifestly have foreseen. A speaker who conjoins the two pieces of information in (5b) must intend an interpretation on which the effort of processing them jointly is justified by extra effects. Such an interpretation would be achieved if, for example, it was relevant to know why Peter took out his key, or how he opened the door.

The criterion of consistency with the principle of relevance provides a means of bridging the gap between what is linguistically encoded and what is ostensively communicated. Of a range of possible hypotheses about the intended interpretation, all of which would yield enough effects to make the utterance worth his attention, the hearer should choose the most accessible one, the one that is easiest to construct. Although other hypotheses might yield adequate effects, this is the only one to yield adequate effects *for the minimum justifiable effort*, and thus satisfy the criterion of consistency with the principle of relevance.

So why did Peter take out his key? How did he open the door? Well, we all have an easily accessible encyclopaedic schema for taking out a key and using it to unlock a door. On hearing (5b), it is natural to interpret it in accordance with this schema, as communicating that he used the key to open the door. No other hypothesis comes more readily to mind. If, on this basis, the hearer can achieve an overall interpretation consistent with the principle of relevance, his hypothesis will automatically be confirmed. A maxim of orderliness is neither necessary nor sufficient to account for this interpretation. Similar arguments apply to the other examples in (4) and (5) above, making the invention of further maxims unnecessary.

¹ For further discussion and a range of additional examples, see Carston (1988).

Recently, Regina Blass (1989, 1990) has used the criterion of consistency with the principle of relevance to argue against an encoding account of a rather different type. Sissala, a Niger-Congo language, has two words for ‘and’. These words are intersubstitutable in certain contexts but carry different implications: *a* suggests that the event described in the second conjunct happened in the normal or obvious way, while *ka* suggests that it was somehow special, abnormal or unexpected. Thus, the Sissala equivalent of (6a) would suggest that Peter lit the fire in the normal way – say in the hearth – while (6b) would suggest that either the fact that Peter lit a fire, or the way he lit it, was unexpected (along lines that the context should help to narrow down):

- (6a) Peter entered the room *a* lit a fire.
 (6b) Peter entered the room *ka* lit a fire.

These differences could be accounted for by lexical stipulation: treating *ka* as encoding a conventional implicature of unexpectedness, for example. Blass suggests a more interesting approach.

She notes, first, that (6a) and (6b) are not syntactically equivalent. *Ka* is a sentence conjunction, *a* a VP conjunction: thus (6b) contains an extra phonetically unrealised S node and Subject NP, making it costlier to process. A speaker aiming at optimal relevance, who can achieve her intended effects by use of (6a), should therefore prefer (6a) to (6b). It follows that the only legitimate interpretation of (6b) is one not achievable by use of (6a). What could such an interpretation be?

By the arguments given above for (5b), (6a) should be understood, where possible, in terms of an encyclopaedic schema for entering a room and lighting a fire. In these circumstances, a speaker who intends something other than the interpretation that would be achieved by use of this schema will not be able to achieve it by (6a). Here the costlier (6b) comes into its own as a vehicle for the less stereotypical interpretation. In this way, Blass shows how the differences between (6a) and (6b) can arise without being linguistically encoded.

Her analysis is confirmed by the cancellability test. If an encoding account were correct, conjoined sentences with *ka* should always carry connotations of unexpectedness; on Blass’s relevance-theoretic account, these connotations should only arise where a less costly alternative, such as (6a), is manifestly available. The crucial examples are thus sentences such as (7), where the two conjuncts have different subjects and conjunction-reduction is impossible, so that no manifestly less costly alternative exists:

(7) Today Peter played football *ka* Mary played golf.

Since the Sissala equivalents of (7) need carry no connotations of unexpectedness, the relevance-theoretic analysis is confirmed.²

From the first, Grice's analysis of 'and' ran into a more serious problem, which could not be solved merely by modifying or replacing the maxims: it seemed to undercut the very possibility of a pragmatic account. According to Grice, pragmatic principles make little or no contribution to the truth conditions of utterances. He regarded (4a) and (4b) above as not only semantically but also truth-conditionally equivalent: their temporal and causal connotations were not part of the proposition expressed, but arose only at the level of implicature. But if this is so, as Cohen (1971) points out, the proposition expressed by (8a) is of the form *P or P*, and the utterance should be redundant; and the proposition expressed by (8b) is of the form *P and not P*, and the utterance should be contradictory:

- (8a) I'm not quite sure what happened: either Peter got angry and Mary left, or Mary left and Peter got angry.
- (8b) What happened was not that Peter got angry and Mary left, but that Mary left and Peter got angry.

The fact that these utterances are perfectly acceptable creates a serious problem for Grice's account.

In recent work, Robyn Carston (1988) has shown how to solve the problem and save the pragmatic approach. Grice assumed that the proposition expressed by an utterance is, essentially, recovered by decoding, and that the only contribution made by the maxims was at the level of what was implicated rather than what was said. In *Relevance*, we challenged this assumption. We argued that although the logical form of an utterance is recovered by decoding, its fully propositional form is obtained by inferential enrichment of the linguistically encoded logical form, constrained by the criterion of consistency with the principle of relevance. It is the propositional form of an utterance, not its logical form, that determines the proposition

² It might be argued that Grice's maxim of brevity could account for these examples, *ka* being longer than *a*. Such an analysis would be empirically distinguishable from ours. We claim that the pragmatic differences between (6a) and (6b) result not from the fact that *ka* is longer than *a* but from the fact that (6b) contains extra, phonetically unrealised syntactic material. Even if *ka* and *a* were identical in length, (6b) would be costlier to process and thus, on our account but not on Grice's, should still have the implications described.

expressed. Carston has shown that Grice's problems disappear if the temporal and causal connotations of utterances such as (4a) and (4b) are treated not as implicatures, but as pragmatically determined aspects of the proposition expressed, which contribute to truth conditions and fall under the scope of logical operators and connectives.³ Her analysis thus confirms the view that the inferential phase of comprehension is not restricted to the recovery of implicatures. We will return to this point below.

5. Conceptual and non-conceptual encoding

The distinction between conceptual and non-conceptual encoding has been explored in recent work by Diane Blakemore (see Blakemore 1987, 1988, 1992; see also Blass 1990; Gutt 1991; Moeschler 1989a,b; Luscher 1989). The idea behind it is this. Inferential comprehension involves the construction and manipulation of conceptual representations; linguistic decoding feeds inferential comprehension; linguistic constructions might therefore be expected to encode two basic types of information: concepts or conceptual representations on the one hand, and procedures for manipulating them on the other.

In the course of comprehension, an utterance is assigned a series of representations, phonetic, phonological, syntactic and conceptual. A *conceptual* representation differs from a phonetic, phonological or syntactic representation in two main respects. First, it has logical properties: it enters into entailment or contradiction relations, and can act as the input to logical inference rules. Second, it has truth-conditional properties: it can describe or partially characterise a certain state of affairs.

Consider (9):

(9) Peter told Mary that he was tired.

Let us suppose that the logical form of (9) looks something like (10a), which is completed into the fully propositional form (10b) by an inferential process of reference assignment:

(10a) x told y at t_1 that z was tired at t_2 .

(10b) Peter Brown told Mary Green at 3.00 p.m. on June 23 1992, that Peter Brown was tired at 3.00 p.m. on June 23 1992.

³ For discussion of Carston's proposals, see Récanati (1989),

Then both the logical form (10a) and the fully propositional form (10b) are conceptual representations, the first recovered purely by decoding and the second by a combination of decoding and inference. The higher-level explicatures derived by embedding (10b) under various propositional-attitude or speech-act descriptions are further examples of conceptual representations recovered from (9) by a combination of decoding and inference.

The idea that there are expressions whose function is not so much to encode a concept as to indicate how to 'take' the sentence or phrase in which they occur has played an important role in pragmatics: in particular, in the work of Ducrot and his associates (Ducrot 1972, 1973, 1984; Anscombe and Ducrot 1983). In speech-act theory, such expressions are treated as illocutionary-force indicators; in the Gricean framework, they are treated as carrying conventional implicatures (for discussion of Grice's treatment, see below).

Within relevance theory, the idea that an expression may encode procedural constraints on the inferential phase of comprehension was first put forward by Diane Blakemore (see Brockway 1981, Blakemore 1987, 1988, 1992). Consider (11), which we have divided into sub-parts (a) and (b):

(11) (a) Peter's not stupid. (b) He can find his own way home.

This utterance has two possible interpretations, which would be encouraged, respectively, by the formulations in (12a) and (12b):

(12a) Peter's not stupid; so he can find his own way home.

(12b) Peter's not stupid; after all, he can find his own way home.

On the first interpretation, (11a) provides evidence for a conclusion drawn in (11b); on the second, (11a) is confirmed by evidence provided in (11b). Blakemore argues that discourse connectives such as 'so' and 'after all' should not be seen as encoding concepts. They do not contribute to the truth conditions of utterances, but constrain the inferential phase of comprehension by indicating the type of inference process that the hearer is expected to go through. As Blakemore points out, such expressions contribute to relevance by guiding the hearer towards the intended contextual effects, hence reducing the overall effort required.

In terms of the distinctions drawn in section 1, Blakemore's semantic constraints on relevance are both procedural and non-truth-conditional. On her approach, 'so' and 'after all' do not encode concepts, and do not contribute to the truth conditions of utterances; instead, they guide the

inferential phase of comprehension. Blakemore's analysis of discourse connectives raises an interesting theoretical question: are the truth-conditional and the conceptual, the non-truth-conditional and the procedural necessarily linked? Does the fact that an expression is truth-conditional entail that it encodes a concept, and the fact that an expression is procedural entail that it encodes a procedure? In later sections, we will argue that it does not. In the next section we will compare Blakemore's account of discourse connectives with Grice's.

6. Explicit and implicit conceptual encoding

Blakemore's work on discourse connectives amounts to a reanalysis in procedural terms of Grice's notion of conventional implicature. Grice does not talk in terms of a conceptual/procedural distinction. Nonetheless, he seems to have thought of the conventional implicatures carried by discourse connectives such as 'but', 'moreover', 'so' and 'on the other hand' in conceptual rather than procedural terms. For one thing, his choice of the term 'implicature' suggests that he thought of conventional implicatures, like conversational implicatures, as distinct propositions with their own truth conditions and truth values. Moreover, he talks in almost identical terms of what was conventionally implicated and what was said, noting, for instance, that items or situations are 'picked out by', or 'fall under', both what was conventionally implicated and what was said.

The difference between conventional and conversational implicatures was, of course, that conventional implicatures were semantically decoded, whereas conversational implicatures were not decoded but inferred. The difference between saying and conventionally implicating was that the truth conditions of the utterance were determined by what was said, whereas conventional implicatures were non-truth-conditional. In terms of the distinctions drawn in section 1, then, Grice appears to treat conventional implicatures as linguistically encoded conceptual representations which make no contribution to the truth conditions of the utterances which carry them, but contribute rather to implicatures. His analysis shows how a linguistic expression which is non-truth-conditional might nonetheless encode conceptual rather than procedural information.

At various points in his writings, Grice analyses 'but', 'moreover', 'on the other hand' and 'so' in terms of his notion of conventional implicature. To illustrate his approach, we will look at his treatment of 'on the other hand' in the 'Retrospective Epilogue' (Grice 1989: 362). Consider (13):

- (13) My brother-in-law lives on a peak in Darien; his great aunt, on the other hand, was a nurse in World War I.

Grice saw the speaker of (13) as asserting that her brother-in-law lived on a peak in Darien and that his great aunt was a nurse in World War I, and conventionally implicating that she has in mind some contrast between these two assertions:

What was asserted by (13):

- (a) The speaker's brother-in-law lives on a peak in Darien.
- (b) The brother-in-law's great aunt was a nurse in World War I.

What was conventionally implicated by (13):

- (a) and (b) contrast in some way.

Grice seems to have thought of conventional implicatures in standard speech-act terms, as indicating the type of speech act performed. Thus, he says of (13):

'Speakers may be at one and the same time engaged in performing speech acts at different but related levels. One part of what [the speaker of (13)] is doing is making what might be called ground floor statements about the brother-in-law and the great aunt, but at the same time as he is performing these speech acts he is also performing a higher-order speech act of commenting in a certain way on the lower-order speech acts. He is *contrasting* in some way the performance of some of these lower-order speech-acts with others, and he signals his performance of this higher-order speech act in his use of the embedded enclitic phrase "on the other hand". The truth or falsity ... of his words is determined by the relation of his ground-floor speech acts to the world; consequently, while a certain kind of misperformance of the higher-order speech-act may constitute a semantic offense, it will not touch the truth-value ... of the speaker's words.' (p. 362)

Notice here the striking similarity between Grice's talk of 'higher-order speech acts' performed by discourse connectives and the relevance-theoretic notion of a 'higher-level explicature'. This raises an interesting question about where the borderline between explicit and implicit communication should be drawn. Grice, like Blakemore, treats the discourse connectives as contributing to implicit rather than explicit communication. Roughly speaking, he equates what is explicitly communicated with what is said (i.e. truth-conditional content), so that all non-truth-conditional constructions are automatically seen as falling on the implicit side.

We do not follow him on this. In *Relevance* (1986: 182) we offered a definition of explicitness and degrees of explicitness:

Explicitness:

An assumption communicated by an utterance U is *explicit* if and only if it is a development of a logical form encoded by U.

On the analogy of ‘implicature’, we call an explicitly communicated assumption an *explicature*. Logical forms are ‘developed’ into explicatures by inferential enrichment. Every explicature, then, is recovered by a combination of decoding and inference, and the greater the element of decoding, the more explicit it will be.

As noted above, our category of explicatures includes not only the proposition expressed by the utterance, but a range of higher-level explicatures obtained by embedding the proposition expressed under an appropriate speech-act or propositional-attitude description. If Grice is right to claim that discourse connectives convey linguistically encoded information about ‘higher-order speech acts’, they would in our framework be analysed as contributing to explicit rather than implicit communication. In general, relevance theorists see the explicit side of communication as much richer, and involving a much greater element of pragmatic inference, than Gricean pragmatists have thought.

Leaving this issue aside for the moment, let us return to semantics proper, and consider whether discourse connectives such as ‘so’, ‘after all’, ‘on the other hand’, etc., are best analysed in conceptual or procedural terms. Grice’s conceptual analysis can be directly compared with Blakemore’s, since both offer analyses of ‘so’. Consider (14):

(14) (a) It’s raining. (b) So the grass is wet.

According to Grice, the use of ‘so’ in (14) indicates that the speaker is ‘performing the speech-act of explaining’, with (14a) being put forward as an explanation of (14b):

What was said by (14):

- (a) It’s raining.
- (b) The grass is wet.

What was conventionally implicated by use of ‘so’:

- (a) explains (b).

According to Blakemore, ‘so’ is an inferential connective indicating that the assumption which follows it is a conclusion. On her account, (14b) is put forward as a conclusion drawn from (14a):

Propositions expressed by (14)

- (a) It's raining.
- (b) The grass is wet.

Procedural information encoded by 'so':

Process (14b) as a conclusion.

Notice first that there are purely descriptive reasons for preferring Blakemore's account: Grice's analysis does not work for all uses of 'so'. (15) is one of Blakemore's examples. The speaker sees someone arrive home laden with parcels and says:

- (15) So you've spent all your money.

Here, there is no explanatory clause corresponding to (14a). The speaker is not *explaining* the fact that the hearer has spent all her money, but drawing a conclusion from an observation she has made. Blakemore's account fits (15) better than Grice's.

In fact there are uses of 'so' which look like counterexamples to any speech-act account. Consider (16a), understood as an indirect speech report of (16b):

- (16a) Peter thought that Mary had a holiday, so he should have one too.
- (16b) Peter thinks, 'Mary had a holiday, so I should have one too'.

(16a) is compatible with Blakemore's inferential account. Though not drawing an inference herself, the speaker of (16a) is attributing a certain inference to Peter. By contrast, she is neither performing a speech act of explanation herself, nor attributing any speech act to Peter: she is reporting thoughts, not words. This suggests that what is needed is not a better speech-act analysis of 'so', but a cognitive analysis such as the one Blakemore has proposed.

Leaving this objection aside, let us recast Grice's analysis so that it avoids the descriptive problem in (15). This could be done by treating (17a) as encoding the conventional implicature in (17b):

- (17a) P, so Q.
- (17b) Q is a consequence of P.

This modified Gricean account is directly comparable with Blakemore's: the only difference between them is that one is conceptual and the other is procedural. Is there any way of choosing between the two accounts?

There is one piece of direct evidence in favour of Blakemore's approach and against the Gricean treatment. Most 'conventional implicatures' are carried by so-called discourse connectives: 'so', 'now', 'well', 'moreover', 'however', and so on. Discourse connectives are notoriously hard to pin down in conceptual terms. If 'now' or 'well' encodes a proposition, why can it not be brought to consciousness? Why is it so hard for non-native speakers of German to grasp the meaning of 'ja' and 'doch'? How can the results of Ducrot's complex analyses of 'but' and other connectives be at once so simple and so insightful? The procedural account suggests an answer to these questions. Conceptual representations can be brought to consciousness: procedures can not. We have direct access neither to grammatical computations nor to the inferential computations used in comprehension. A procedural analysis of discourse connectives would explain our lack of direct access to the information they encode.

There are two further types of construction whose analysis provides indirect evidence for Blakemore's procedural account of discourse connectives and against a Gricean conceptual account. In the next section, we will look at some non-truth-conditional expressions which, unlike the discourse connectives, clearly call for conceptual treatment. In the following section, we will look at some non-truth-conditional constructions which clearly call for procedural treatment. Indirect evidence for Blakemore's account of discourse connectives is that they seem to have more in common with constructions in the procedural than the conceptual class.

7. Proposition expressed versus higher-level explicatures

In section 3, we distinguished the proposition expressed by an utterance from its higher-level explicatures. In section 5, we argued that from a cognitive point of view, these higher-level explicatures are conceptual representations, capable of entailing and contradicting each other and representing determinate states of affairs. Though true or false in their own right, they do not generally contribute to the truth conditions of their associated utterances. Mary's utterance in (1b) above is true or false depending on whether she can or can't help Peter find a job, not on whether she does or doesn't say, or believe, or regret that she can't help him.

Now consider the utterances in (18):

(18a) Seriously, I can't help you.

- (18b) Frankly, I can't help you.
- (18c) Confidentially, I can't help you.
- (18d) Unfortunately, I can't help you.

Illocutionary adverbials such as 'seriously', 'frankly' and 'confidentially', and attitudinal adverbials such as 'unfortunately', are standardly treated as making no contribution to the truth conditions of utterances in which they occur. Récanati says of the attitudinal adverb 'happily':

'Deleting the adverb would not change the proposition expressed by the sentence ... because the modification introduced by the adverb is external to the proposition and concerns the speaker's emotional attitude to the latter. This attitude is neither "stated" nor "described", but only "indicated"' (Recanati 1987: 50)

Here we will consider only illocutionary adverbials, and we will take for granted their non-truth-conditional status (for more detailed discussion, see Ifantidou, this volume). The main point we want to make is that, even though illocutionary adverbials are clearly non-truth-conditional, there are good reasons to treat them as encoding concepts.

Notice, first, that even if the illocutionary adverbials in (18) are non-truth-conditional, their synonymous manner-adverbial counterparts in (19) must clearly be treated as encoding concepts which contribute to the truth conditions of the associated utterances in the regular way:

- (19a) Mary told Peter seriously that she couldn't help him.
- (19b) Mary said frankly to Peter that she couldn't help him.
- (19c) Mary informed Peter confidentially that she couldn't help him.

Given this, the simplest hypothesis is that in (18) they encode exactly the same concepts. The only difference is that in interpreting (18), the hearer must incorporate these concepts into a higher-level explicature some elements of which are not encoded but inferred. The fact that the illocutionary adverbials make no contribution to the truth conditions of (18) would then follow from the more general fact that the higher-level explicatures with which they are associated make no contribution to truth conditions either. This analysis fits well with standard speech-act accounts of illocutionary adverbials, on which an illocutionary adverb such as 'seriously' is seen not as contributing to the proposition expressed by the utterance, but as modifying the type of speech-act performed (see for example Bach and Harnish 1979, chapter 10, section 3; Nolke 1990).

By contrast, a procedural analysis of illocutionary adverbials would run into serious difficulties. Firstly, as has often been pointed out, an utterance like (20) is ambiguous, with the two possible interpretations in (21):

- (20) Seriously, are you leaving?
 (21a) I ask you seriously whether you are leaving.
 (21b) I ask you to tell me seriously whether you are leaving.

This is not surprising on the explicature account. Whenever (20) is interpretable as a request to tell, the illocutionary adverb should be interpretable as modifying either the requesting or the telling. It is not obvious how this ambiguity could be handled in procedural terms.

Secondly, many sentence adverbials are semantically complex. Consider (22a–d):

- (22a) Frankly speaking, he has negative charisma.
 (22b) Speaking frankly, though not as frankly as I'd like to, he isn't much good.
 (22c) In total, absolute confidence, how are you getting on with Maria?
 (22d) While he's out getting the coffee, what did you think of Bill's talk?

Such compositionality is unsurprising if illocutionary adverbials encode conceptual representations, which can undergo semantic interpretation rules in the regular way. It is not obvious what compositionality would mean in procedural terms.

Thirdly, in some cases at least, the speaker who uses an illocutionary adverbial can lay herself open to charges of untruthfulness in its use. Consider (23)–(25):

- (23a) *Mary*: Frankly, this steak is less than perfect.
 (23b) *Peter*: That's not true. You're not being frank.
 (24a) *Mary*: Seriously, what a gorgeous tie.
 (24b) *Peter*: That's not true. You're never serious.
 (25a) *Mary*: Now I've brought you your fourth whisky, what did you think of the play?
 (25b) *Peter*: That's not true. It's only my third.

If illocutionary adverbials encode elements of conceptual representations which can be true or false in their own right, such exchanges are not surprising.

In fact, in some cases an illocutionary adverbial seems to contribute directly to the truth conditions of the associated utterance. Consider (26):

(26a) *Peter*: What can I tell our readers about your private life?

(26b) *Mary*: On the record, I'm happily married; off the record, I'm about to divorce.

If the illocutionary adverbials 'on the record' and 'off the record' made no contribution to the truth conditions of (26b), then Mary's utterance should be perceived as contradictory; yet intuitively it is not. But if these adverbials contribute to truth conditions, then a fortiori they encode conceptual representations, and the procedural analysis is disconfirmed.

It seems, then, that there is good reason to treat illocutionary adverbials as both non-truth-conditional and conceptual, thus abandoning the idea that all non-truth-conditional meaning is necessarily procedural and cut to a single pattern.

8. Constraints on explicatures and constraints on implicatures

We have now illustrated three of the four logically possible types of meaning distinguished in section 1:

- (a) Most regular 'content' words, including the manner adverbials 'seriously', 'frankly', etc., are *conceptual and truth-conditional*: they encode concepts which are constituents of the proposition expressed by the utterance, and hence contribute to the utterance's truth conditions.
- (b) Various types of sentence adverbial, including the illocutionary adverbials 'seriously', 'frankly', etc., are *conceptual and non-truth-conditional*: they encode concepts which are constituents not of the proposition expressed but of higher-level explicatures.
- (c) Discourse connectives such as 'so' and 'after all' are *procedural and non-truth-conditional*: they encode procedural constraints on implicatures.

In this section, we will argue that personal pronouns such as 'I' and 'you' illustrate the fourth category of meaning: they are both *procedural and truth-conditional*.

The idea that there are procedural constraints on truth-conditional content was suggested (in different terms) by Jakobson and Benveniste in their discussion of 'shifters'. However, when Benveniste (1966: 252) says that the pronoun 'I' means 'the speaker of the utterance in which the token of "I"

occurs', his proposal is seriously ambiguous. Kaplan (1989) points out (again in different terms) that the claim that 'I' means 'the speaker' has different consequences depending on whether it is conceptually or procedurally understood.

Suppose that David Kaplan says (27):

(27) I do not exist.

Then if 'I' is treated as encoding the *concept the speaker*, (27) will express the proposition in (28):

(28) The speaker of (27) does not exist.

But if 'I' is treated merely as encoding an instruction to identify its referent by first identifying the speaker, then (27) will express the proposition in (29):

(29) David Kaplan does not exist.

These two propositions differ in their truth conditions. (29) is true in any state of affairs in which David Kaplan does not exist. (28) is true in any state of affairs in which (27) is uttered and its speaker does not exist. Since such a state of affairs is impossible, if (27) expressed the proposition in (28), it would be necessarily false. Kaplan argues that though (27) is false whenever it is uttered, it is not *necessarily* false. The proposition it expresses is true in any state of affairs in which David Kaplan does not exist. In other words, (27) must be understood as expressing (29), not (28).

Accordingly, Kaplan proposes to distinguish the *content* of an expression from its *character*. The content of 'I' in (27) is the individual David Kaplan; the character of 'I' is a rule for identifying its content in any given context. Such rules, Kaplan comments,

'tell us for any possible occurrence of the indexical what the referent would be, but they do *not* constitute the content of such an occurrence. Indexicals are directly referential. The rules tell us what it is that is referred to. Thus, they *determine* the content (the propositional constituent) for a particular occurrence of an indexical. But they are not a *part* of the content (they constitute no part of the propositional constituent).' (Kaplan 1989: 523)

In terms of the distinctions drawn in section 1, this amounts to the claim that 'I' and other pronouns are both truth-conditional and procedural, thus

illustrating the fourth logically possible type of encoded meaning, and refuting the assumption that there is a necessary linkage between the truth-conditional and the conceptual.⁴

We have now looked at two quite different types of procedural expression: discourse connectives and pronouns. Both constrain the inferential phase of comprehension by reducing the hypothesis space that has to be searched in arriving at the intended interpretation. Discourse connectives impose constraints on *implicatures*: they guide the search for intended contexts and contextual effects. Pronouns impose constraints on *explicatures*: they guide the search for the intended referent, which is part of the proposition expressed. This raises the possibility that there might be a still further type of procedural expression, which constrains not the proposition expressed by an utterance but its higher-level explicatures.

At the end of *Relevance*, we drew attention to a range of constructions which seem to us to be best analysed in these terms. The idea that declarative sentences and their non-declarative counterparts express the same propositions but perform different speech acts is familiar from speech-act theory. While there are serious problems with the speech-act approach to non-declarative sentences (for detailed discussion, see Wilson and Sperber 1988), we believe that the semantic differences between declarative sentences and their non-declarative counterparts can be successfully analysed as differences not in the propositions they express but in the higher-level explicatures they communicate: for example, a declarative utterance should be treated as a case of *saying that*, and an imperative utterance as a case of *telling to*.⁵ Notice that this proposal, like the one for 'I' above, can be understood in two different ways. On one interpretation, Mary's utterance in (1b) above would be treated as conceptually encoding the higher-level explicature 'the speaker says that she can't help Peter':

(1a) *Peter*: Can you help me?

(1b) *Mary* (sadly): I can't.

Understood in this way, our proposal would be a variant of the performative

⁴ Within a relevance-theoretic framework, Ruth Kempson has been developing a procedural approach to anaphora in interesting recent work (see Kempson 1988; see also Ariel 1990, Kleiber 1990, Reboul 1990). The analysis of pronouns thus looks like providing an important source of evidence about the nature of procedural constraints on interpretation.

⁵ In a suitably attenuated sense on which to *say that* P, for example, is to make no commitment to the truth of P.

hypothesis abandoned for excellent reasons many years ago (on the history of the performative hypothesis, see Levinson 1983). On the other interpretation - the one proposed in *Relevance* - what is encoded is not a conceptual representation but a set of hints for constructing one. The content of this higher-level representation will be partially determined by contextual information, and will specify the illocutionary force of the utterance in terms of much richer concepts than the abstractions 'saying that' or 'telling to'. As we said in *Relevance*:

'illocutionary force indicators such as declarative or imperative mood or interrogative word order merely have to make manifest a rather abstract property of the speaker's informative intention: the direction in which relevance is to be sought.' (Sperber and Wilson 1986: 254)

That is, illocutionary-force indicators should be seen as encoding procedural constraints on the inferential construction of higher-level explicatures. It seems clear that this interpretation is to be preferred. (For details of this approach to non-declaratives, see Wilson and Sperber 1988; Clark forthcoming.)

As is well known, the functions performed in English by mood and word order are performed in many other languages by so-called discourse or illocutionary particles. Certain dialects of French, for example, have an interrogative particle *ti*, which appears to perform the same functions as word-order inversion does in other dialects. If word-order inversion is correctly analysed as encoding not a concept but a constraint on higher-level explicatures, then by the same arguments, illocutionary particles such as *ti* (at least such particles as are fully integrated into the syntax, i.e. are genuine parts of the language) should be analysed in similar terms. Perhaps the question particle 'eh' in English might be a candidate for similar treatment.

In the framework of relevance theory, Regina Blass (1990) has analysed the 'hearsay' particle *re* in Sissala as encoding a constraint on explicatures. Perhaps some 'attitudinal' discourse particles (again, to the extent that they are fully integrated into the language) might be analysed on similar lines. When Mary uses the dissociative particle 'huh!' in (30), for example, she might be seen as encouraging the construction of the higher-level explicature in (31):

(30) Peter's a genius, huh!

(31) Mary doesn't think that Peter's a genius.

Within this category of procedural constraints on explicatures, there is thus a rich variety of data to explore.

For discourse particles such as ‘*ti*’, the failure of the performative hypothesis provides direct evidence against an analysis in terms of conceptual encoding and for a procedural account. Returning to the analysis of discourse connectives such as ‘*so*’ and ‘*after all*’, their obvious similarities to discourse particles provide indirect evidence against an account in terms of conceptual encoding and for a procedural account.

In this section, we have proposed that certain pronouns, illocutionary-force indicators and discourse particles should all be analysed as encoding procedural constraints on explicatures. The pronouns are truth-conditional and contribute to the proposition expressed; illocutionary-force indicators and discourse particles are non-truth-conditional and contribute to higher-level explicatures. These differences between them should not, we feel, be allowed to obscure the important similarities between the types of meaning they encode.

9. Conclusion

In section 1, we sketched two contrasting approaches to linguistic semantics: one focusing on utterances, their truth conditions and the speech acts they are used to perform, the other locating utterances within a broader cognitive framework. Throughout this paper we have taken a resolutely cognitive approach. We assume, in fact, that the primary bearers of truth conditions are not utterances but conceptual representations; to the extent that utterances have truth conditions, we see these as inherited from the propositions those utterances express. We have tried to show that an approach along these lines can yield genuine insight into the varieties of linguistically encoded meaning.

What, then, of the more traditional linguistic approach? Surely there is still a consistent, coherent set of intuitions about the truth conditions of utterances which any adequate theory of linguistic semantics needs to explain? There may be, though we know of no systematic effort to show that this is so. Most semanticists simply assume (as we have throughout this paper) that every utterance encodes a single logical form, expresses a single proposition and has a single set of truth conditions. We do not know how far this assumption can be maintained. In utterances with illocutionary adverbials, or parentheticals of the type discussed by Ifantidou (1992), Itani (1990)

and Blakemore (1991), it might be argued, à la Grice, that the speaker is simultaneously making two assertions, each with its own truth conditions; one might then investigate the possibility that intuitions about the truth conditions of the utterance as a whole are based on the assertion which makes the major contribution to overall relevance. Clearly, much research remains to be done in this area. What we hope to have shown is that such research can be usefully conducted within the broader cognitive and communicative framework outlined here.

References

- Anscombe, J. and O. Ducrot, 1983. *L'argumentation dans la langue*. Brussels: Madarga.
- Ariel, M., 1990. *Accessing noun phrase antecedents*. London: Routledge.
- Bach, K. and R. Harnish, 1979. *Linguistic communication and speech acts*. Cambridge, MA: MIT Press.
- Benveniste, E., 1966. *Problèmes de linguistique générale*. Paris: Gallimard.
- Blakemore, Diane, 1987. *Semantic constraints on relevance*. Oxford: Blackwell.
- Blakemore, Diane, 1988. *So as a constraint on relevance*. In: R. Kempson (ed.), *Mental representation: The interface between language and reality*, 183–195. Cambridge: Cambridge University Press.
- Blakemore, Diane, 1991. *Performatives and parentheticals*. *Proceedings of the Aristotelian Society* 91(3), 197–213.
- Blakemore, Diane, 1992. *Understanding utterances*. Oxford: Blackwell.
- Blass, Regina, 1989. *Pragmatic effects of co-ordination: The case of 'and' in Sissala*. *UCL Working Papers in Linguistics* 1: 32–51. University College London.
- Blass, Regina, 1990. *Relevance relations in discourse: A study with special reference to Sissala*. Cambridge: Cambridge University Press.
- Brockway, Diane, 1981. *Semantic constraints on relevance*. In: H. Parret, M. Sbisà and J. Verschueren (eds.), *Possibilities and limitations of pragmatics*, 57–78. Amsterdam: Benjamins.
- Carston, Robyn, 1988. *Implicature, explicature and truth-theoretic semantics*. In: R. Kempson (ed.), *Mental representation: The interface between language and reality*, 155–181. Cambridge: Cambridge University Press. Reprinted in S. Davis (ed.), *Pragmatics: A reader*, 33–51. Oxford: Oxford University Press.
- Clark, B., forthcoming. *A relevance-theoretic analysis of pseudo-imperatives*. To appear in *Linguistics and Philosophy*.
- Cohen, L.J., 1971. *Some remarks about Grice's views about the logical particles of natural languages*. In: Y. Bar-Hillel (ed.), *Pragmatics of natural languages*, 50–68. Dordrecht: Reidel.
- Ducrot, O., 1972. *Dire et ne pas dire*. Paris: Hermann.
- Ducrot, O., 1973. *La preuve et le dire*. Paris: Mame.
- Ducrot, O., 1980. *Analyses pragmatiques*. *Communications* 32, 11–60.
- Ducrot, O., 1984. *Le dire et le dit*. Paris: Minuit.
- Fodor, J., 1983. *The modularity of mind*. Cambridge, MA: MIT Press.

- Grice, H. P., 1989. *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Gutt, Ernst-August, 1991. *Translation and relevance*. Oxford: Blackwell.
- Ifantidou, E., 1992. *Parentheticals and relevance*. Paper delivered to the LAGB Conference, Brighton.
- Itani, R., 1990. *Explicature and explicit attitude*. UCL Working Papers in Linguistics 2, 52–64.
- Kaplan, David, 1989. *Demonstratives*. In: J. Almog, J. Perry, H. Wettstein (eds.), *Themes from Kaplan*, 481–563. Oxford: Oxford University Press.
- Kempson, Ruth, 1988. *On the grammar–cognition interface: The principle of full interpretation*. In: R. Kempson (ed.), *Mental representations: The interface between language and reality*, 199–224. Cambridge: Cambridge University Press.
- Kleiber, Georges, 1990. *Marqueurs référentiels et processus interprétatifs: Pour une approche 'plus sémantique'*. Cahiers de linguistique française 11.
- Levinson, S., 1983. *Pragmatics*. Cambridge: Cambridge University Press.
- Luscher, Jean-Marc, 1989. *Connecteurs et marques de pertinence: L'exemple de d'ailleurs*. Cahiers de linguistique française 10.
- Moeschler, J., 1989a. *Modélisation du dialogue: Représentation de l'inférence argumentative*. Paris: Hermes.
- Moeschler, J., 1989b. *Pragmatic connectives, argumentative coherence and relevance*. *Argumentation* 3.
- Nolke, H., 1990. *Pertinence et modalisateurs d'énonciation*. Cahiers de linguistique française 11.
- Reboul, A., 1990. *Rhétorique de l'anaphore*. In: G. Kleiber, J. Tyvaert (eds.), *L'anaphore et ses domaines*, 279–300. Paris: Klincksieck.
- Récánati, F., 1987. *Meaning and force*. Cambridge: Cambridge University Press.
- Récánati, F., 1989. *The pragmatics of what is said*. *Mind and Language* 4(4), 295–329.
- Sperber, Dan and Deirdre Wilson, 1986. *Relevance: Communication and cognition*. Oxford: Blackwell, and Cambridge, MA: Harvard University Press.
- Wilson, Deirdre and Dan Sperber, 1988. *Mood and the analysis of non-declarative sentences*. In: J. Dancy, J. Moravcsik, C. Taylor (eds.), *Human agency: Language, duty and value*, 77–101. Stanford, CA: Stanford University Press.