

LINGUISTIC FORM AND RELEVANCE

Deirdre Wilson and Dan Sperber

Abstract

In this paper, we survey the ways in which information may be conveyed by an utterance. We consider a variety of theoretically relevant distinctions: between conveying and communicating, encoding and inferring, linguistic and non-linguistic communication, saying and indicating, explicating and implicating, conceptual and procedural, truth-conditional and non-truth-conditional information. Our main focus is on a range of constructions which have been treated in the literature as encoding non-truth-conditional information: we argue that not all such constructions function in the same way. In particular, we distinguish constructions encoding constraints on implicatures (cf. Blakemore 1987) from those encoding constraints on explicatures, a category hitherto unrecognised.

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*, *descriptive* 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 Recanati 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 about the various speech acts it is intended to perform.

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

The main point we want to make in this paper is that these two responses are not equivalent. They do not classify the data in identical ways. It is true that many linguistic constructions which are truth-conditional on the first approach are representational on the second, and many constructions which are non-truth-conditional on the first approach are procedural on the second. However, all four possibilities are realised. Some constructions are truth-conditional on the first approach but procedural on the second, and some constructions are non-truth-conditional on the first approach but representational on the second. This raises a general question. Are both these distinctions really necessary? Is the four-way classification theoretically justified? Or can one of them be dispensed with?

These internal questions about the decoding phase of comprehension are mainly of interest to linguistic semanticists. 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 and 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 (see following page) we have tried to draw the threads together and give a general picture of the various types of information that an utterance might convey. 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 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. 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 information-transmission is that the hearer has special help in recognising the speaker's informative intention. Ostensive communication communicates 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. That 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):

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

Suppose that in saying (1b), Mary expected Peter not only to notice her sadness, 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 have those in (3):

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

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:

- (4) a. Peter got angry and Mary left.
- b. Mary left and Peter got angry.

The temporal implications 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 implications 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 implications: (4a) suggests that Mary left because Peter got angry and (4b) suggests the reverse. These implications do not follow from a maxim of orderliness alone. Or consider (5a-d):

- (5) a. Peter went into the kitchen and found Mary.
- b. Peter took out his key and opened the door.
- c. Mary injured her leg and sued Peter.
- d. 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 implications is linguistically encoded: all are cancellable without contradiction. The problem they raise 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 implications 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: that is, an interpretation which yields adequate effects for the minimum justifiable processing effort in a way the speaker could manifestly have foreseen. Since the contextual effects of (5b) will vary depending on how Peter opened the door - did he open it with his key, for instance, or batter it down with an axe? - in normal circumstances the speaker cannot be supposed to have been indifferent as to which of these assumptions is made.

The criterion of consistency with the principle of relevance provides a method for resolving indeterminacies as to the intended interpretation. Of a range of possible hypotheses, all of which would yield enough effects to make the utterance worth his attention, the hearer should choose the most obvious one, the one that is easiest to make. 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 how did Peter 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. 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 about (5b) will be confirmed, and the speaker will be taken to have licensed the assumption that immediately after taking out his key, Peter used it to unlock the door. 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.

Recently, Regina Blass (Blass 1989, forthcoming) 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

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

'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 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 something unexpected - say, that he set fire to the curtains:

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

One way of accounting for these differences would be by lexical stipulation. 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 convey it by means of (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 implications of unexpectedness; on Blass's relevance-theoretic account, these should only arise where a less costly alternative, such as (6a), is available. The crucial examples are thus sentences such as (7), where the two conjuncts have different subjects and conjunction-reduction is impossible:

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

The Sissala equivalents of (7) need carry no implications of unexpectedness, and 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 reformulating 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 an

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.

utterance. He regarded (4a) and (4b) above as not only semantically but also truth-conditionally equivalent: their temporal and causal implications were not part of the proposition expressed, but arose only at the level of implicature. But if this is so, as Cohen (1971) pointed out, the proposition expressed by (8a) is of the form *P or P*, and the utterance should be tautologous; and the proposition expressed by (8b) is of the form *P and not P*, and the utterance should be contradictory:

- (8) a. I'm not quite sure what happened: either Peter got angry and Mary left, or Mary left and Peter got angry.
- b. 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. It is the propositional form of an utterance, not its logical form, that determines the proposition expressed. Carston has shown that Grice's problems disappear if the temporal and causal implications 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; Blass forthcoming; Gutt forthcoming; Moeschler 1989a, 1989b; Luscher 1989). The idea behind it is this. Inferential comprehension involves the construction and manipulation of conceptual representations. Linguistic decoding feeds inferential comprehension. An utterance might therefore be expected to encode two basic types of information: conceptual representations on the one hand, and information about how to manipulate them on the other.

In the course of comprehension, an utterance is assigned a series of representations. A *conceptual* representation differs from a phonetic 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.

³ For discussion of Carston's proposals, see Recanati 1989.

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 reference assignment:

- (10) a. x told y at t1 that z was tired at t2.
 b. Peter Brown told Mary Green at 3.00 p.m. on June 23, 1989,
 that Peter Brown was tired at 3.00 p.m. on June 23, 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.

Within relevance theory, the idea that an utterance might encode constraints on the inferential phase of comprehension was first put forward by Diane Blakemore (see Brockway 1981; Blakemore 1987). Consider (11):

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

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

- (12) a. Peter's not stupid; so he can find his own way home.
 b. 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 '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 make an obvious contribution to relevance by cutting down the hearer's processing costs.

In terms of the distinctions drawn in section 1, Blakemore's semantic constraints on relevance are both procedural and non-truth-conditional. The information they encode can be seen, on the one hand, as guiding the inferential phase of comprehension, and on the other as indicating the type of speech act the speaker intended to perform. In the next section, we will look at an alternative account of these phenomena; in later sections, we will consider some cases where the two distinctions begin to come apart.

6. Explicit and implicit conceptual encoding

Blakemore's work amounts to a reanalysis in cognitive terms of Grice's notion of conventional implicature. 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.

In his 'Retrospective Epilogue' (Grice 1989), Grice saw the speaker of (13) as conventionally implicating that she has in mind some contrast between the states of affairs described in the first and second conjuncts of her utterance. An implicature, by definition, makes no contribution to the truth conditions of the utterance which conveys it. However, it seems clear that Grice regards conventional implicatures as conceptual representations with logical properties, capable in their own right of being true or false: he talks, for example, (p. 360) of items or situations as being 'picked out' by, or 'falling under' both what was conventionally implicated and what was said. In terms of the distinctions drawn in section 1, Grice's conventional implicatures are linguistically encoded conceptual representations which make no contribution to the truth conditions of the utterances which carry them. His analysis thus shows how information that is non-truth-conditional on the first approach might be representational on the second.

Grice himself sees conventional implicatures in standard speech-act terms. 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)

Grice's proposals can be directly compared with Blakemore's, since they both offer analyses of 'so'. According to Grice, the speaker of (14) is 'performing the speech-act of *explaining*' (p. 362):

- (14) a. The sun's shining. b. So I'm happy.

On his account, (14a) is put forward as an explanation of (14b). 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).

There are reasons for preferring Blakemore's account to Grice's. In the first place, 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):

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

It might be possible to save part of Grice's analysis by recasting it in cognitive terms but keeping the assumption that 'so' encodes a conceptual representation rather than a procedure. For example, 'so' in (14) might be treated as encoding one of the propositions in (17):

- (17) a. The fact that the speaker of (14) is happy is a conclusion.
- b. (14b) was derived by inference.

Is there any way of choosing between these two versions of the cognitive account?

There is one direct argument that favours Blakemore's version. Most 'conventional implicatures' are carried by so-called discourse particles: 'so', 'now', 'well', 'moreover', 'however', and so on. Discourse particles are notoriously hard to pin down in conceptual terms. If 'now' or 'well' encodes a proposition, why can this proposition not be brought to consciousness? Why is it so hard for non-native speakers of German to grasp the meaning of 'ja' and 'doch'? The procedural account suggests an answer to this question. 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 particles 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 a procedural account of discourse particles and against a conceptual account. In the next section, we will look at some expressions which are much better candidates for conceptual treatment than those Grice had in mind. In the following section, we will look at some constructions which clearly call for procedural treatment. Indirect evidence for Blakemore's account of discourse particles is that they seem to have more in common with constructions in the procedural than the representational class.

7. Proposition expressed versus higher-level explicatures

In section 3, we distinguished the propositional form of 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 the utterances which carry them. 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):

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

Sentence adverbs such as 'seriously', 'frankly' and 'confidentially' are often treated as making no contribution to the truth conditions of utterances in which they occur. Recanati says of the sentence 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)

Nolke (forthcoming) treats the class of adverbs he calls 'adverbes de l'enonciation' in similar terms.

Whatever the merits of these intuitions, it is clear that in utterances such as (19) these adverbs must be treated as encoding concepts which contribute to truth conditions in the regular way:

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

Our 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 these sentence 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 Grice's idea that the speaker may on the one hand perform a 'ground floor' speech act, and on the other perform a 'higher-order speech act of commenting in a certain way on the lower-order speech act'. For these examples, Grice's proposal seems to us to be correct.

By contrast, a procedural analysis of sentence adverbials would run into serious difficulties. Firstly, as Nolke (forthcoming) points out, an utterance like (20) is ambiguous, with the two possible interpretations in (21):

(20) Seriously, are you leaving?

- (21) a. I ask you seriously whether you are leaving.
b. 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 sentence 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 highly complex. Consider (22a-d):

- (22) a. Frankly speaking, he has negative charisma.
b. Speaking frankly, though not as frankly as I'd like to, he isn't much good.
c. Since we're alone for the first time since you had that argument, how are you getting on with Maria?
d. While he's out getting the coffee, what did you think of Bill's talk?

The resulting interpretations are straightforwardly dealt with on the assumption that sentence adverbials encode conceptual representations, which are typically compositional. It is not obvious what compositionality would mean in procedural terms.

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

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

If sentence adverbials encode elements of conceptual representations which, though not contributing to the truth conditions of the utterance which conveys them, can be true or false in their own right, such exchanges are not surprising.

In fact, in some cases a sentence adverbial does seem to contribute to the truth conditions of the utterance which conveys it. Consider (26):

- (26) a. Peter: What can I tell our readers about your private life?
b. Mary: On the record, I'm happily married; off the record, I'm about to divorce.

If the sentence 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 the truth conditions of the utterances which convey them, then a fortiori they encode conceptual representations, and the procedural analysis is

disconfirmed.

8. Constraints on explicatures and constraints on implicatures

So far, we have found evidence for the following types of encoded information:

- (a) logical forms, i.e. conceptual information which contributes both to explicit content and to truth conditions;
- (b) constraints on implicatures, i.e. procedural information which contributes neither to explicit content nor to truth conditions;
- (c) constituents of higher-level explicatures, i.e. conceptual information which contributes to explicatures but not, by hypothesis, to truth conditions.

We have found no evidence for the implicit counterpart to (c):

- (d) constituents of implicatures, i.e. conceptual information which contributes to implicatures but not to truth conditions.

In this section, we will argue that there is one further category, the explicit counterpart to Blakemore's semantic constraints on relevance:

- (e) constraints on explicatures, i.e. procedural information which contributes to explicit content and, on occasion, to truth conditions.

The idea that there are procedural constraints on truth-conditional content has already been explored in the literature on formal semantics. It is often claimed that the pronoun 'I' refers to the speaker. Kaplan (1989) points out that this claim 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)

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

At the end of *Relevance*, we pointed out a second source of evidence: the analysis of non-declarative sentences. We proposed that declaratives and their imperative counterparts, for instance, express the same proposition but differ in their higher-level explicatures: a declarative utterance is a case of saying that; an imperative utterance is a case of *telling to*.⁴ This proposal, like the meaning rule for 'I' above, can be taken in two ways. On one interpretation, Mary's utterance in (1b) above would conceptually encode the higher-level description 'x says that...'. On this interpretation, our proposal would be a variant of the notorious performative hypothesis abandoned for excellent reasons many years ago (on the history of the performative hypothesis, see Levinson 1983). On the other interpretation, what is encoded is not a conceptual representation but a set of hints for constructing one. As we put it 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 the relevance of the utterance is to be sought.' (Sperber and Wilson 1986: 254)

It seems clear that this interpretation is to be preferred.⁵

A third source of evidence is a class of particles with illocutionary or attitudinal functions. Certain dialects of French, for example, have a question particle 'ti', which achieves the same effects as word-order inversion does in other dialects. If we are right, 'ti' encodes not a conceptual representation but a constraint on higher-level explicatures; the question particle 'eh?' in English might be a candidate for similar treatment.

In the framework of relevance theory, Regina Blass (forthcoming) has discussed the 'hearsay' particle 're' in Sissala in terms which suggest that it

⁴ In a suitably attenuated sense on which to say that P, for example, is to make no commitment to the truth of P.

⁵ For a development of this approach to non-declaratives, see Wilson and Sperber 1987.

should be analysed as encoding a constraint on explicatures. On such an analysis, the Sissala equivalent of (30) would carry the explicature in (31):

- (30) It'll be fine today, *re*.
(31) According to *x*, it'll be fine today.

Certain attitudinal indicators might be analysed on similar lines. When Mary uses the dissociative particle 'huh!' in (32), for example, she might be seen as encouraging the construction of the explicature in (33):

- (32) Peter's a genius, huh!
(33) Mary doesn't think that Peter's a genius.

For illocutionary and attitudinal particles, the failure of the performative hypothesis provides direct evidence against conceptual encoding and for a procedural account. Returning to the analysis of discourse particles, their obvious similarities to illocutionary particles provide indirect evidence against an account in terms of conceptual encoding and for a procedural account.

In this section, we have proposed that indexicals, speech-act indicators and attitudinal indicators should all be analysed as encoding constraints on explicatures. One important difference between them should, however, be noted. Indexicals contribute to the proposition expressed by an utterance; speech-act and attitudinal indicators contribute to its higher-level explicatures: as a result, they contribute to truth conditions in rather different ways.

The proposition expressed by an utterance has a certain truth-conditional content; the higher-level explicatures determine how this proposition is to be taken. Indexicals, by determining what proposition is expressed, make an obvious contribution to truth conditions, but explicatures make a contribution too. To know whether an utterance even *has* truth conditions, one has to know whether it is a case of *saying*, *telling*, or *asking*. To know what those truth conditions are, one has to know whether the proposition expressed is used echoically, as in (30) and (32), or descriptively, as in (1b). If we are right, then not only indexicals but all the constructions discussed in this section belong to the final category of encoded information discussed in section 7:

- (e) constraints on explicatures, i.e. procedural information which contributes to explicit content and, on occasion, to truth conditions.

9. Conclusion

For purposes of this paper, we have assumed that every utterance encodes a single logical form, expresses a single proposition and has a single set of truth conditions. However, we are not sure how far this assumption can be maintained. In utterances with sentence adverbials, or with parentheticals of the type discussed by Blakemore (forthcoming), it might be argued, *à la* Grice, that the speaker is simultaneously producing two utterances, each with its own truth conditions. With illocutionary and attitudinal particles such as 'eh?', 'huh!' and 're', a similar possibility might be investigated. It is not obvious, then, that there is a consistent and comprehensive enough set of intuitions about *the* truth conditions of an utterance to bear much theoretical weight.

What would be lost if the linguistic distinction between truth-conditional and non-truth-conditional meaning collapsed? We would be left with the pair of distinctions around which this paper was constructed: the cognitive distinction between conceptual and procedural, and the pragmatic distinction between explicit and implicit. As we have seen, the conceptual-procedural distinction provides its own partition of phenomena into truth-conditional and non-truth-conditional, and captures many of the distinctions which the linguistic distinction was set up to explain. What remains is the class of higher-level explicatures some element of which, apart from the proposition expressed, is linguistically encoded. The survival of the linguistic distinction between truth-conditional and non-truth-conditional meaning will depend on there being a consistent, coherent set of intuitions about these.

References

- Blakemore, Diane (1987) *Semantic constraints on relevance*. Blackwell, Oxford.
- Blakemore, Diane (1988) 'So' as a constraint on relevance. In R. Kempson (ed.) *Mental Representation: The language-reality interface*. CUP, Cambridge.
- Blakemore, Diane (forthcoming) Explicit performatives. To appear in *Cahiers de linguistique française* 11.
- Blass, Regina (1989) Pragmatic effects of co-ordination: the case of 'and' in Sissala. *UCL Working Papers in Linguistics* 1.
- Blass, Regina (forthcoming) *Relevance relations in discourse: A study with special reference to Sissala*. To be published by CUP, Cambridge.
- Brockway, Diane (1981) Semantic constraints on relevance. In H. Parret, M. Sbisà and J. Verschueren (eds) *Possibilities and limitations of pragmatics*. J. Benjamin, Amsterdam.
- Carston, Robyn (1988) Implicature, explicature and truth-theoretic semantics. In R. Kempson (ed.) *Mental representation: The language-reality interface*. CUP, Cambridge.
- Cohen, L. J. (1971) Some remarks on Grice's views about the logical particles of natural languages. In Y. Bar-Hillel (ed.) *Pragmatics of natural languages*. Reidel, Dordrecht.
- Fodor, Jerry (1983) *The modularity of mind*. MIT Press, Cambridge, Mass.
- Grice, H. P. (1989) *Studies in the way of words*. Harvard UP, Cambridge, Mass.
- Gut, Ernst-August (forthcoming) *Translation and relevance*. To be published by Blackwell, Oxford.
- Kaplan, David (1989) Demonstratives. In J. Almog, J. Perry and H. Wettstein (eds) *Themes from Kaplan*. OUP, Oxford.
- Kempson, Ruth (forthcoming) *Language and cognition: licensing grammar*. To be published by Blackwell, Oxford.
- Kleiber, Georges (forthcoming) Marqueurs référentiels et processus interprétatifs: pour une approche 'plus sémantique'. To appear in *Cahiers de Linguistique française* 11.
- Levinson, S. (1983) *Pragmatics*. CUP, Cambridge.
- Luscher, Jean-Marc (1989) Connecteurs et marques de pertinence: l'exemple de d'ailleurs. *Cahiers de Linguistique française* 10.

- Moeschler, Jacques (1989a) *Modélisation du dialogue: Représentation de l'inférence argumentative*. Hermes, Paris.
- Moeschler, Jacques (1989b) Pragmatic connectives, argumentative coherence and relevance. *Argumentation* 3.
- Nolke, H. (forthcoming) Pertinence et modalisateurs d'énonciation. To appear in *Cahiers de Linguistique Française* 11.
- Reboul, Anne (forthcoming) Rhétorique de l'anaphore.
- Récanati, François (1981) *Les énoncés performatifs*. Minuit, Paris.
- Récanati, François (1987) *Meaning and Force*. CUP. Cambridge.
- Récanati, François (1989) The pragmatics of what is said. *Mind and Language* 4,4.
- Sperber, Dan and Deirdre Wilson (1986) *Relevance: Communication and cognition*. Blackwell, Oxford.
- Wilson, Deirdre and Dan Sperber (1988) Mood and the analysis of non-declarative sentences. In J. Dancy, J. Moravcsik and C. Taylor (eds) *Human Agency: Language, Duty and Value*. Stanford UP, Stanford.




Figure 1