UCL Working Papers in Linguistics 10 (1998)

Recent approaches to bridging: Truth, coherence, relevance

DEIRDRE WILSON AND TOMOKO MATSUI

Abstract

This paper considers three recent approaches to bridging reference, based on notions of truth, coherence and relevance, and argues that a relevance-based approach to bridging is preferable on both descriptive and explanatory grounds. Using questionnaire results from Matsui 1995, it compares the predictions of the relevance-theoretic account with several versions of truth-based and coherence-based approaches, and shows how the intuitions behind these approaches might be explained on a relevance-theoretic account.

1 Introduction

Some classic examples of bridging reference are given in (1):

(1) a. Jane has a new house. *The front door* is blue.
b. Susan sent her meal back. *The steak* was underdone.

In interpreting the italicised NPs *the front door*, *the steak*, the hearer has to supply a referent which is not explicitly mentioned in the preceding discourse, but whose existence can be inferred from what has gone before. One way of doing this would be by adding the following assumptions to the context in which the italicised NPs are interpreted:

(2) a. Jane's new house has a front door.
b. Susan's meal included steak.

Such assumptions may be seen as 'bridging the gap' between what has been explicitly stated in the first part of the utterance and what is needed for reference assignment in the second. The need for such assumptions is what mainly distinguishes bridging

reference from other types of reference. We will call them *bridging assumptions*, and the italicised NPs *bridging NPs*.

Early analyses of bridging (e.g. Clark & Haviland 1977, Sanford & Garrod 1981) were conducted within a largely Gricean pragmatic framework, with its machinery of Co-operative Principle and maxims, maxim-violation, inferential intention-recognition, and so on. Bridging assumptions such as those in (2) were treated as Gricean implicatures: beliefs that had to be attributed to the speaker in order to preserve the assumption that she was obeying the Co-operative Principle and maxims in saying what she said. In developing their account of bridging, Clark and Haviland added further maxims to the framework, stating felicity conditions on the use of specific linguistic constructions, whose violation, real or apparent, would trigger the bridging process. For example, the bridging implicatures in (2) would be triggered by violation of a Maxim of Antecedence governing the use of definite NPs.

Early analyses of bridging thus raised pragmatic questions on two quite different levels. On the one hand, there were *general* pragmatic questions having to do with the justification of an overall pragmatic framework and its consequences for the analysis of various pragmatic processes. On the other, there were *specific* pragmatic questions having to do with the felicity conditions on the use of particular constructions (e.g. definite descriptions, referential expressions), and the particular effects they achieve. In this paper, our concern is mainly with general pragmatic questions. We have focused on the analysis of bridging more for the light it can shed on pragmatic processes in general than for the particular problems raised by the analysis of definite NPs.

Most recent approaches to pragmatics start from the Gricean assumption that the hearer, in interpreting an utterance, is looking not just for some arbitrary interpretation but for the one overtly intended by the speaker. Most of them also share the following assumptions about the general form that pragmatic processes such as disambiguation, reference resolution, identification of implicatures, and so on, should take:

(3) a. Candidate interpretations (e.g. disambiguations, reference assignments, implicatures, and so on) differ in their accessibility to the hearer, and are therefore evaluated in a certain order.
b. They are evaluated in terms of some criterion of pragmatic acceptability that the resulting overall interpretation is supposed to meet.
c. The first interpretation that satisfies the pragmatic criterion is the one the hearer should choose.

Here, assumption (3a) concerns *accessibility*, or *effort*, (3b) concerns *acceptability*, or *effect*, and (3c) concerns the relation between the two. In practice, psycholinguists tend to focus more on questions of effort, and pragmatists (or discourse analysts, or text linguists) on questions of effect. However, an adequate pragmatic framework should integrate the two.

On the effort side, it is clear that a hearer who is looking for the overtly intended interpretation is not always justified in simply choosing the first interpretation that comes to mind; so a pragmatic criterion based on considerations of accessibility alone is unlikely to work. On the effect side, two main types of criterion are currently on offer: (a) those based on Grice's maxims of Quality and/or Quantity, which assume that hearers expect utterances (or discourses) to be true, or true and informative, or true and informative and evidenced; and (b) those based on Grice's maxim of Relation, which assume that hearers also expect utterances (or discourses) to be relevant, or coherent. We will argue that theories of type (a) are insufficient, and that we should be trying to construct theories of type (b). We will then compare two type (b) theories - coherence theory and relevance theory - and try to show that relevance theory offers the best way of integrating not only the effort and effect factors, but also all the intuitions that underlie the truth-based and coherence-based approaches.

2. Pragmatic approaches to bridging

2.0 Let us start with a brief survey of approaches to bridging, organised around the twin factors of accessibility, or effort, and acceptability, or effect. Many of our examples are drawn from a questionnaire on bridging developed by Tomoko Matsui (and discussed in more detail in Matsui 1995). While the numbers involved in this questionnaire were small, the results are quite suggestive, and we will use them to illustrate the type of problems that an adequate analysis of bridging would have to solve.

2.1 Accessibility

In the case of bridging, two main types of factor affect the accessibility of candidate referents:

(4) a. the accessibility of the linguistic antecedent (e.g. *the house* in (1a));
b. the accessibility of the bridging assumption.

We will say a word about each – briefly, since it seems clear that hearers do not simply choose the first interpretation that occurs to them, and it's the pragmatic criterion used in evaluating this interpretation that mainly concerns us here.

2.1.1 Accessibility of linguistic antecedents. In the general literature on reference resolution, a lot of work has been done on the factors affecting accessibility of linguistic antecedents. These include order of mention, syntactic position, recency of mention, manner of mention, thematic role, semantics of the main verb, parallel function, choice of conjunction, and overall salience in discourse (for survey and discussion, see Matsui 1995). The focus-based (or centering) approaches developed by e.g. Grosz (1981), Sidner (1983) and Grosz & Sidner (1986), and adopted both in psycholinguistics (Sanford, Moar & Garrod 1988) and formal semantics (Asher and Wada 1988), have been particularly fruitful: the constituent claimed to be 'in focus' is precisely the most accessible linguistic antecedent for reference assignment.[1]

Most work on bridging assumes that what would be the most accessible antecedent for direct reference would also be the most accessible antecedent for bridging. This is explicitly argued in Erku & Gundel (1987), who consider examples like the following:

(5) We went to a Thai restaurant. *The waitress* was from Bangkok.

(6) We stopped for drinks at the New York Hilton before going to the Thai restaurant. *The waitress* was from Bangkok.

In (5), *the waitress* would be understood as the one in the Thai restaurant. In (6), there are two possible antecedents, *the Hilton* and *the Thai restaurant*, and the question is how *the waitress* in (6) would be understood. According to Erku & Gundel, she would be understood as the one in the Hilton, despite the strong encyclopaedic association between *Thai* and *Bangkok*. They explain this by appeal to the relative accessibility of the two antecedents, as predicted by Sidner's 'Expected Focus Algorithm', which picks out the focus on the basis of syntactic position and thematic role. In (6), the Expected Focus Algorithm would pick out the main clause NP *the Hilton* rather than the subordinate-clause NP *the Thai restaurant*; this focused constituent would then be the expected antecedent for the bridging NP *the waitress*. By the same token, the expected antecedent for *the baby orang-utan* in (7) would be *the Hilton*, not *the zoo*.

[1]For further discussion, see Reboul 1994; Reboul et al. 1997.

(7) ?We stopped for drinks at the Hilton before going to the zoo. *The baby orang-utan* was really cute.

According to Erku & Gundel, (7) would be perceived as unacceptable or stylistically infelicitous because the most accessible antecedent has to be rejected on the basis of encyclopaedic knowledge about where baby orang-utans are to be found. Their analysis thus makes the following two claims: (a) if at all possible, the hearer chooses the most accessible antecedent; and (b) if the most accessible interpretation has to be rejected, as in (7), the utterance will be considered stylistically infelicitous.

Matsui's questionnaire contains some examples designed to test these claims. In each case, the bridging NP is *the humidity*, and there are two possible antecedents: *England* and *Hong Kong*. Subjects were asked to answer the question, 'Where did the humidity not bother John?', and to note any utterance they regarded as stylistically infelicitous. Percentage responses are shown in brackets after each example:

(8) a. John worked in England before moving to Hong Kong 5 years ago. *The humidity* didn't bother him. [40% England; 60% Hong Kong]
b. John worked in Hong Kong before moving to England 5 years ago. *The humidity* didn't bother him. [60% Hong Kong; 40% England]
c. John worked in England before moving to Hong Kong 5 years ago. *The humidity* really bothered him. [40% England; 60% Hong Kong]
d. John worked in Hong Kong before moving to England 5 years ago. *The humidity* really bothered him. [100% Hong Kong]

These responses show that subjects did not always choose the most accessible antecedent as predicted by the Expected Focus Algorithm. Their preferences were split in all cases except (8d), which is the only one to behave as expected. Subjects were also asked to note any utterance they regarded as stylistically infelicitous. Only one person noted an infelicity in (8), and that went counter to Erku & Gundel's predictions: it was in (8b), with the interpretation 'Hong Kong', which according to the Expected Focus Algorithm is the more accessible antecedent.

In fact most work on reference resolution, including Sidner's, acknowledges, first, that the accessibility of referents is affected by a variety of factors other than thematic role; and second, that the most accessible candidate can be rejected and another selected on pragmatic grounds. The differences in interpretation of the examples in (8) provide

obvious confirmation of this point. Although we will continue to use the Expected Focus Algorithm as a rough guide to the accessibility of antecedents, our main concern will be with the pragmatic criterion used to accept or reject interpretations. As we will show, this criterion itself sheds some light on questions of accessibility, since one of the factors affecting the accessibility of interpretations is likely to be the hearer's expectation about the effects the speaker is aiming to achieve.

2.1.2 Accessibility of bridging assumptions. A second factor affecting the accessibility of candidate referents is the accessibility of bridging assumptions such as (2a-b) above. Early experimental work showed that bridging in standard examples like (1a-b) is extremely fast. Sanford & Garrod (1981) explained this in terms of frame-based or scenario-based associations between the linguistic antecedent (e.g. *the house* in (1a)) and the candidate referent. The idea is that processing of the linguistic antecedent in (1a-b) would automatically activate encyclopaedic information containing the frame-based bridging assumptions in (2a-b). A caricature version of this idea is found in Hawkins 1978:

(9) The bridging assumption must belong to a frame associated with the antecedent. Otherwise the utterance will be unacceptable.

Matsui's questionnaire contained a range of examples which on the whole bore out this common-sense assumption that accessibility of encyclopaedic knowledge affects the acceptability of bridging. Subjects were given sets of four examples and asked to rank their acceptability on a 7-point scale (with 7 being the most acceptable and 1 the least). Here is an illustration, with the mean ratings shown in brackets after each example:

(10) a. I was looking at van Gogh's self-portrait. *The missing ear* made me feel sad. [6.54]
b. I went to see some impressionist paintings. *The missing ear* made me feel sad. [4.08]
c. I went to a gallery yesterday. *The missing ear* made me feel sad. [3.23]
d. I spent all day in London. *The missing ear* made me feel sad. [1.46]

The research was done at a time when a van Gogh exhibition in London was receiving a lot of publicity. Even so, the interpretation of each successive example involves a little more creativity (hence, effort), and as predicted on frame-based accounts, the greater the creativity, the lower the acceptability.

However, Erku & Gundel 1987 show that a frame-based approach is inadequate on its own to explain the acceptability or unacceptability of bridging. They compare examples such as (11a-c):

(11) a. ?A car went by. *The dog* was barking.
b. Don't go near that house. *The dog* will bite you.
c. Don't go near the car. *The dog* will bite you.

The frame-based criterion in (9) above would rule out (11a) as unacceptable, because dogs are not part of the frame for cars. However, as Erku & Gundel point out, (11b) would normally be acceptable, even though dogs are not normally part of the frame for houses; moreover, (11c) shows that an acceptable link can be made between car and dog. Indeed, even (11a) would be acceptable in certain circumstances, e.g. as a denial of (12):

(12) Dogs never bark in cars.

Again the conclusion must be that while the accessibility of bridging assumptions is a factor affecting overall pragmatic acceptability, it is not the only one.

The arguments of this section confirm the view that, whatever the accessibility facts, they can be overridden if the resulting interpretation is not pragmatically acceptable. The question is, what constitutes pragmatic acceptability? Psycholinguists have rarely been seriously interested in this question; pragmatic theorists must be.

2.2 Acceptability

2.2.0 According to Grice, hearers can identify the intended interpretation of an utterance because utterances raise certain expectations. The Gricean expectation is, of course, that the speaker was obeying the Co-operative Principle and maxims (or at least the Co-operative Principle) in saying what she said. Most recent approaches to pragmatics have started out as attempts to develop one or more of Grice's maxims into an adequate pragmatic criterion. We will look first at type (a) theories, based on the maxims of Quality and/or Quantity (truthfulness and/or informativeness) and then at type (b) theories, based on the maxim of Relation, 'Be relevant'.

2.2.1 Truth-based approaches. What we are calling truth-based approaches are built around the Quality and/or Quantity maxims. A prototype example is Lewis (1983), who

suggests that the hearers should use something like the following pragmatic criterion in interpreting bridging cases:

(13) Accept the first candidate referent that leads to an overall interpretation that is true, informative and evidenced. (Lewis 1983: 242)

On Lewis's view, reference is assigned to the most salient (or accessible) candidate; if the resulting interpretation is pragmatically unacceptable, saliencies are reordered according to the following rule, until an acceptable interpretation is found:

(14) *Rule of accommodation for comparative salience*
If at time *t* something is said that requires, if it is to be acceptable, that *x* be more salient than *y*; and if, just before *t*, *x* is no more salient than *y*; then - ceteris paribus and within certain limits - at *t*, *x* becomes more salient than *y*. (Lewis 1983: 242)

In this case, the motivation for the reinterpretation process is the expectation that the utterance will be true, informative and evidenced.

Some of Matsui's questionnaire examples were designed to test the claim that the most accessible candidate may be rejected in favour of a less accessible one which leads to a more factually plausible interpretation. Here are some cases which seem to fit this view:

(15) a. I prefer England to Italy. I hate *the pasta* there. [100% Italy]
b. I prefer the country to the town. *The traffic* really bothers me. [100% town]
c. I prefer London to Edinburgh. I hate *the snowy winters*. [100% Edinburgh]
d. I prefer Italy to England. *The weather* is worse. [60% Italy; 40% England]

In these examples, the most accessible antecedent (according to the Expected Focus Algorithm) is the direct object of *prefer*, but this is rejected in favour of a less accessible candidate, and it could be argued that the reason is that the resulting interpretation is more likely to be true.

The question is how this argument could be worked out in detail. Several families of truth-based criteria could be proposed, each making different predictions. First, a truth-based criterion might be strong or weak. A strong criterion would compare competing interpretations and choose the most factually plausible (or plausible and informative) one; a weak criterion would accept the first interpretation that reaches a certain threshold of factual plausibility (or plausibility and informativeness, etc.), without necessarily comparing it with alternative interpretations. Second, a truth-based criterion may be local or global. A local criterion would pick out (for example) the most factually plausible

interpretation of the utterance currently being processed; a global criterion would aim to maximise the factual plausibility of the discourse as a whole. Finally, the predictions of a criterion based on truth alone would differ from those of a criterion based on truth and informativeness, truth and evidencedness, or truth, informativeness and evidencedness. Without considering all these possibilities, we would like to suggest that none of them is likely to work unless some appeal to a notion of coherence or relevance is made.

Consider, first, a strong, local criterion, which compares alternative interpretations of the utterance currently being processed and chooses the one that is most likely to be true (with a minimal check for consistency with prior discourse, but no expectation of relevance or coherence). One consequence of such a strong criterion is that the accessibility factor in comprehension would have to be discounted: the fact that one interpretation was more accessible than others would no longer be any reason for accepting it. Another consequence is that all possible interpretations would have to be computed and compared, which in many cases would lead to a combinatorial explosion. Moreover, some of the questionnaire results do not seem to fit this criterion. Consider (16):

(16) I prefer the restaurant on the corner to the student canteen. *The cappuccino* is less expensive.

The most accessible antecedent for *the cappuccino* (according to the Expected Focus Algorithm) is the direct object *the restaurant.* However, it is part of encyclopaedic knowledge that cappuccino in a student canteen is likely to be less expensive than cappuccino in an ordinary restaurant. If the hearer simply chooses the most factually plausible interpretation of the utterance currently being processed, we might therefore expect the most accessible candidate to be rejected, and *the cappuccino* in (16) to be understood as referring to the cappuccino in the student canteen. In the questionnaire, this hypothesis was tested by asking the following questions:

(17) a. Where is the cappuccino less expensive? [Restaurant: 100%]
b. In general, where do you get less expensive cappuccino, a student canteen or a restaurant? [Canteen: 100%]

Here, the results were unequivocal: 100% of subjects said that *the cappuccino* in (16) referred to the cappuccino in the restaurant; and 100% of subjects said that the cappuccino in a student canteen is generally less expensive than the cappuccino in a restaurant. So the hypothesis that the hearer chooses the most factually plausible interpretation of the utterance currently being processed is not borne out.

Consider now a weak local criterion, which simply accepts (again with a minimal check for consistency with prior discourse but no expectation of coherence or relevance) the first interpretation to reach a certain threshold of factual plausibility, which, as in this example, may be rather low. Weak criteria are preferable to strong criteria from the processing-effort point of view. However, while this criterion would explain the results in (16), it no longer explains the interpretation of the examples in (15). In (15a), for instance, the most accessible interpretation of the second part of the utterance (according to the Expected Focus Algorithm) is the one on which the speaker is saying she hates the pasta in England. This interpretation is both factually plausible and consistent with the prior discourse, and the revised pragmatic criterion no longer explains why it is rejected when a much less plausible interpretation of (16) is retained.

Consider now a strong global criterion designed to pick out the interpretation on which the discourse as a whole is most likely to be true (or true and informative, or true, informative and evidenced). It might be thought that this would explain the results in both (15) and (16). For example, it might be suggested that in (15a), the speaker's claim that she prefers England to Italy somehow makes it more plausible that she hates the pasta in Italy; and in (16), her claim that she prefers the restaurant on the corner somehow makes it more plausible that the cappuccino there is less expensive. But - even leaving aside the processing-effort disadvantages of strong criteria - it is not easy to see how this idea could be worked out. After all, it is quite possible to like a country without liking its food, or to like a restaurant despite the fact that it serves expensive coffee. Moreover, someone who says she prefers England to Italy would not normally be taken to implicate that she prefers the food in England; and someone who says she prefers one restaurant to another would not normally be taken to implicate that the coffee in her preferred restaurant is less expensive. So it is hard to see how the first part of these utterances could be seen as providing evidence for the second.

We all feel, of course, that the only *pragmatically* plausible interpretation of (15a) is the one on which the speaker prefers England to Italy *because* she hates the pasta in Italy, and that the only pragmatically plausible interpretation of (16) is the one on which the speaker prefers the restaurant on the corner *because* it serves less expensive coffee. It might be argued that these interpretations can be dealt with in a truth-based framework by treating the second part of these utterances as providing evidence for the first. But there are several problems with this approach.

One problem is that hearers may be able to choose between two interpretations, both of which provide evidence for an earlier statement. Consider (18):

(18) I ran from the classroom to the playground. *The children* were making too much noise.

This utterance has two interpretations: the speaker may have run out of the classroom to escape the noisy children; or she may have run into the playground to rebuke the noisy children. A hearer assigning reference according to the strong truth-based criterion would now have to consider not only the intrinsic plausibility of the two possible interpretations of the second part of the utterance, but also the amount of evidence that each of them would provide for the first part of the utterance (or perhaps for the implicitly communicated information that the speaker ran from one place to the other because the children there were making too much noise). More generally, every possible combination of disambiguations, reference assignments and enrichments would have to be compared, not only for its own intrinsic plausibility but for the amount of evidence it provides for earlier statements (or implicatures) in the discourse. The resulting combinatorial explosion would dramatically increase the processing effort required. For empirical arguments against this version of the criterion, see section 4.

Matsui's questionnaire contains another type of example which is hard to explain in terms of this approach. These examples were in fact designed to test the hypothesis that encyclopaedic associations between the bridging NP and one of two roughly equally accessible antecedents might affect the hearer's choice. The idea was that a hearer motivated by considerations of factual plausibility should assign reference to the more strongly associated antecedent. Consider (19):

(19) a. Kevin moved from England to New Zealand. He loves *the sheep*. [100% New Zealand]
b. Kevin moved from New Zealand to England. He loves *the sheep*. [100% England]
c. Kevin moved from England to New Zealand. He hates *the sheep*. [100% New Zealand]
d. Kevin moved from New Zealand to England. He hates *the sheep*. [60% England; 40% New Zealand]

According to Expected Focus Algorithm, the prepositional phrases *from England* and *to New Zealand* are roughly equally accessible as antecedents for the bridging NP *the sheep*, with a marginal preference, if anything, for the goal phrase (e.g. *to New Zealand* in (19a)). At the same time, for British subjects, at least, there is a strong association between New Zealand and sheep, and on this basis one might expect a hearer governed by local

expectations of truth to interpret *the sheep* in all four examples as being in New Zealand. As the questionnaire results show, this expectation is not borne out in the case of (19b), where 100% of subjects preferred the antecedent *England*, or (19d), where preferences were split.

These results present further problems for the strong global criterion envisaged above. Consider (19a) and (19b), with the interpretations shown. How were these interpretations arrived at? According to the strong global criterion, they would be arrived at by treating the second part of the utterance as providing evidence for the statement made in the first part. On this approach, (19a) would be understood as saying that Kevin moved from England to New Zealand *because* he loved the sheep in New Zealand, and (19b) would be understood as saying that Kevin moved from New Zealand to England *because* he loved the sheep in England. This may be so, but there is another equally plausible interpretation on which the second part of the utterance describes not the reason for Kevin's move, but merely a result. In these examples, it is not easy to be sure which interpretations subjects were getting, since they converge on the same antecedent. However, in (19c) they come apart, and it is clear that the Result rather than the Reason interpretation was preferred. This presents a further problem for the strong global criterion, since the Result interpretation does not stand in any obvious evidential relation to the first part of the utterance. To put it more generally, utterances in discourse may be related in various ways. According to the truth-based approach, all these relations must be analysable in terms of the notions of truth, informativeness and evidence. The above examples suggest that things are not as simple as that.

We have attributed to Lewis the view that a pragmatically acceptable interpretation must be true, informative and evidenced. In fact, he adds that other considerations, of relevance, for example, may be involved. The results discussed in this section tend to confirm the view that, as far as pragmatic acceptability is concerned, being true, informative and evidenced is not enough. We now turn to two rival attempts to supply the missing ingredient: coherence-based and relevance-based approaches.

2.2.2 Coherence based approaches. Perhaps the most popular current view in both the psycholinguistic and the computational literature is that the missing ingredient is one of discourse coherence. A prototype example is Hobbs (1979), who says, essentially:

(20) Accept the first candidate referent that yields an overall interpretation which is coherent with prior discourse. (Hobbs 1979)

Here the challenge is to develop a theoretically adequate notion of coherence. Attempts

have been made to analyse this in terms of theoretical notions such as discourse topic or discourse plan (e.g. Reinhart 1981, Giora 1998, Grosz & Sidner 1985) However, by far the most popular view, in both the computational and the psycholinguistic literature, is that coherence can be defined in terms of a set of coherence relations that hold between successive utterances in a discourse: e.g. Elaboration, Temporal Sequence, Cause-Consequence, and so on. It is this relation-based notion of coherence that we will look at here.[2]

The assumption behind coherence-based approaches to comprehension is that hearers expect utterances to be coherent with prior discourse; when an utterance has two possible interpretations, they will choose the one that satisfies their expectation of coherence. Here is a classic example from Jerry Hobbs, showing how the search for coherence can guide reference resolution:

(21) John can open Bill's safe. *He* knows the combination.

According to Hobbs, *he* should be interpreted as referring to John rather than Bill because, on this interpretation, the second part of the utterance would count as an Elaboration (or Explanation) of the first. Approaches along these lines have been developed by Hobbs 1979, Mann & Thompson 1988, Samet & Schank 1984, Asher & Lascarides 1993, etc.

The coherence-based approach seems to shed some light on our questionnaire examples. Returning to the examples in (15), where subjects unanimously rejected the most accessible antecedent in favour of an alternative, less accessible one, it might be argued that they were guided by considerations of coherence. In (15a), for example, by taking *Italy* as the antecedent for *the pasta*, they could interpret the second part of the utterance as giving a Reason, or an Explanation, for why the speaker prefers England to Italy. Similarly, in (16), it might be argued that hearers chose *the restaurant* as antecedent for *the cappuccino* because the resulting interpretation would provide a Reason, or Explanation, for why the speaker preferred the restaurant, and that this was enough to outweigh considerations of factual plausibility.

Although it seems intuitively obvious that something like this is going on, there are many descriptive and theoretical problems with the coherence-based approach. Here, we will mention just two. In the first place, the hearer may be able to choose between two interpretations, both of which would make the discourse coherent. Consider (22):

[2]For further discussion of topic-based approaches to coherence, see Wilson 1998.

(22) a. Jane moved from California to Manchester. She hates *the warm winters*. [100% California]
b. Sara left Australia for England. She hates *the sandy beaches*. [100% Australia]

Each of these examples has two possible interpretations, each involving a different coherence relation. In (22a), for instance, Jane may have moved from California to Manchester *because* she hates the warm winters in California (Explanation); or she may have moved from California to Manchester and *as a result* hate the warm winters in Manchester (Result). Both interpretations are causal: in the Result case the causal relation runs forward (from the first part of the utterance to the second), and in the Explanation case, it runs backwards (from the second part of the utterance to the first). As the questionnaire results show, subjects unanimously chose the Explanation interpretation, in which the causal relation runs backwards. The same point applies to (22b), where subjects again chose the Explanation interpretation. The question is why this is so.

The obvious answer is that coherence relations may themselves be more or less accessible. One attempt to deal with the accessibility of coherence relations in a psychologically motivated way is by Sanders, Spooren and Noordman (1992, 1993). They analyse the coherence relations in terms of a set of features, e.g. Causal versus Additive, Basic versus Non-basic Order, which may be used to predict their accessibility. In fact, the proposed analysis makes the wrong predictions about the interpretation of (22a) and (22b). According to Sanders, Spooren & Noordman, relations involving Basic (forwards) Order should be more accessible than those involving Non-Basic (backwards) Order, and on this approach, interpretations based on Result rather than Explanation should be preferred. More generally, this approach predicts that the accessibility of coherence relations depends on their intrinsic features, so that whatever relation is most accessible in one situation will be most accessible in all others. Yet it is clear that the preferred interpretations of subjects vary from situation to situation. For example, in (19c), as noticed above, 100% of subjects favoured a Result interpretation, even though an Explanation interpretation would have been quite plausible. This suggests that the accessibility of coherence relations cannot be accounted for solely in terms of intrinsic features, as Sanders, Spooren & Noordman imply.

A second problem for coherence theories is that an utterance may have two alternative interpretations, both of which satisfy the *same* coherence relation. An example was given in (18) above:

(18) I ran from the classroom to the playground. *The children* were making too much noise.

Here, as noted above, the bridging NP *the children* has two possible antecedents. Whichever is chosen, the second part of the utterance will provide an Explanation of the event described in the first part. In this case, knowing the intended coherence relation will not solve the problem of reference resolution; yet it is easy to see that there are circumstances in which one or other interpretation would be preferred.

More generally, coherence-based approaches raise descriptive questions about which coherence relations exist, and explanatory questions about why they exist and how they are acquired. We will return to these questions after sketching an alternative attempt to supply the missing ingredient, this time by developing a notion of relevance.

2.2.3 Relevance based approaches. Erku and Gundel (1987) argue that in the interpretation of bridging examples, considerations of relevance are decisive:

> What we would like to suggest here is that the only thing which makes it possible for [bridging] expressions to be recognised as anaphoric is the maxim of relation, i.e. the expectation that the speech act performed in the use of some sentence be relevant to the context in which it occurs (Erku & Gundel 1987: 542).

In their view, what makes the difference between the acceptable and unacceptable examples in (11) above is not discourse coherence but relevance. Similarly, they argue that the stylistic unacceptability of example (7) above is due to some connection between relevance and topic, or expected focus. Their criterion for pragmatic interpretation would therefore be something like the following:

(23) Accept the first candidate referent that yields an overall interpretation that is relevant. (Erku & Gundel 1987)

However, they make no attempt to provide a theoretical definition of relevance, commenting merely that:

> The importance of this notion [relevance] in the pragmatics of natural language has been noted in a number of works... There has, however, been relatively little progress in making the notion explicit. (Erku & Gundel 1987: 543).

For several years, relevance theorists have been trying to develop a notion of relevance

which might form the basis of a psychologically plausible account of comprehension. This is what we will turn to now.

3 An outline of relevance theory

Relevance theory shares the Gricean assumption that hearers are looking for the overtly intended interpretation of an utterance. It differs from the Gricean approach in two main respects. First, it is not maxim-based: it contains no general communicative principles that speakers and hearers have to know and use. Second, it does not assume that communication is necessarily co-operative in Grice's sense: that speaker and hearer have to share a common purpose over and above that of understanding and being understood. Its basic claim is that what is fundamental to communication - because it is fundamental to cognition - is the pursuit of relevance.

On the cognitive level, relevance theory claims that human attention and processing resources go to information that seems relevant. This is expressed as the First, or Cognitive, Principle of Relevance.

(24) *Cognitive Principle of Relevance*
Human cognition tends to be geared to the maximisation of relevance.

Relevance is defined in terms of cognitive effects and processing effort:

(25) *Relevance*
a. The greater the cognitive effects, the greater the relevance;
b. The smaller the effort needed to achieve those effects, the greater the relevance.

Cognitive effects are achieved when new information interacts with existing contextual assumptions in one of three ways:

(26) *Cognitive effects*
a. Strengthening an existing assumption;
b. Contradicting and eliminating an existing assumption;
c. Combining with an existing assumption to yield contextual implications.

Processing effort is affected by two main factors: the form in which the information is presented, and the accessibility of the context. The accessibility factors discussed above

in section 2 would fit in here.

To illustrate these ideas, suppose that Peter wants to get to Boston by plane as soon as possible, and doesn't know when the next plane is. He asks Mary, who may tell him one of three things:

(27) a. The next plane to Boston is at 5.30.
b. The next plane to Boston is sometime after 4.00.
c. The next plane to Boston leaves 7,500 seconds after 3.25.

Which information would be most relevant to Peter? Answer: (27a). This is more relevant than (27b) for reasons of effect: it implies everything that (27b) does, and more besides. It is also more relevant than (27c), this time for reasons of effort. Since 7,500 seconds after 3.25 is in fact 5.30, (27c) has the same cognitive effects as (27a), but more processing effort is needed to derive them.

Moving now to the level of communication, relevance theory proposes a Second, or Communicative, Principle of Relevance, which is justified as follows. To communicate with someone is to offer them information. Offers raise expectations. If I offer you food, you expect it to be edible. If human cognition is relevance-oriented, then offers of information should raise expectations of relevance; and by standard Gricean arguments, if there is an interpretation that the speaker might reasonably have intended to satisfy these expectations, that is the one the addressee should choose.

We define a notion of optimal relevance which is meant to spell out what the addressee is entitled to look for in terms of effort and effect:

(28) *Optimal relevance*
An utterance is optimally relevant to an addressee iff:
a. it is relevant enough to be worth the addressee's processing effort;
b. it is the most relevant one compatible with the speaker's abilities and preferences.

This entitlement is spelled out in the Second, or Communicative, Principle of Relevance:

(29) *Communicative Principle of Relevance*
Every utterance communicates a presumption of its own optimal relevance.

The Communicative Principle of Relevance and the definition of optimal relevance in turn suggest a comprehension procedure which we claim is spontaneously followed in

utterance interpretation:

(30) *Relevance-theoretic comprehension procedure*
   a. consider cognitive effects in their order of accessibility (i.e. follow a path of least effort);
   b. stop when the expected level of relevance is achieved.

This procedure integrates effort and effect in the following way. It claims that the hearer is entitled to expect at least enough cognitive effects to make the utterance worth his attention, that the processing effort is the effort needed to achieve these effects, and that the hearer is entitled to accept the first interpretation that satisfies his expectation of relevance.

To illustrate how these ideas might apply to a straightforward case of bridging, consider one of the questionnaire examples which received a unanimous response:

(22) b. Sara left Australia for England. She hates the sandy beaches. [100% Australia]

Here, the comprehension procedure works smoothly. On the effort side, *Australia* is the most accessible antecedent: first, because it is picked out by the Expected Focus algorithm as the direct object of *left*; and second, because of the encyclopaedic association between *Australia* and *sandy beaches*. Moreover, the second part of the utterance, on this interpretation, could be expected to achieve at least adequate cognitive effects by answering a question implicitly raised by the interpretation of the first part: why did Sara leave Australia for England? To say that the first part of the utterance implicitly raises a question amounts to no more than saying that it provides easy access to a context in which an answer to this question would be relevant enough to be worth the hearer's attention. There is no need to appeal to a coherence relation of Explanation. From the relevance-theoretic point of view, computation of coherence relations is a waste of effort unless it gives rise to otherwise inaccessible effects. The expectation of optimal relevance is enough on its own to account for the interpretation of this example.

One advantage of the relevance-based approach is that it allows for greater flexibility in the relations between utterances in discourse. There is no need to assume a finite set of discrete coherence relations which must be recognised as part of the comprehension process. Part of the appeal of the coherence-based approach rests on its use of relations such as Explanation, Temporal Sequence, Cause-Consequence and so on, which everyone would intuitively recognise as holding between utterances in discourse. But theoretical accounts of coherence relations have repeatedly encountered two problems. First, there are

questions about how fine-grained the relations should be. The number of relations proposed has varied between 3 and 300, with no obvious way of deciding between them (see e.g. Hovy 1990). There are also questions about how many of the relations a pair of utterances can enter into simultaneously. The more fine-grained the relations, and the more of them can be simultaneously recognised, the greater the flexibility in interpretation. But by analysing comprehension in terms of cognitive effects, relevance theory achieves the same flexibility with no additional machinery, and no awkward questions about where to draw the line.

The second problem is that many pairs of utterances do not stand in any of the intuitively obvious coherence relations, and to account for their interpretation in coherence-based terms, apparently arbitrary new relations have to be invented. Consider (33), for instance, another example on which subjects responded unanimously:

(33) I prefer the town to the country. *The traffic* doesn't bother me. [town 100%]

Here the most accessible antecedent is *the town*, picked out by both the Expected Focus Algorithm and the encyclopaedic association between *town* and *traffic*. On this interpretation, the second part of the utterance might be expected to achieve at least adequate cognitive effects by answering questions of the following sort, in a context easily accessible from the interpretation of the first part:

(34) a. Isn't the town too noisy?
b. Wouldn't she rather have the quiet of the country?
c. Isn't there anything she dislikes about the town?

Yet none of the intuitively obvious coherence relations is involved. One could, of course, invent a relation such as 'Anticipated Objection' and claim that it must be identified as part of the process of identifying the intended effects. But this would add a pointless extra element of effort, which could be eliminated by bypassing the coherence relation and going directly to the effects. In this way, relevance theory accounts for the intuitions behind the coherence-based approach, but without the machinery of coherence relations.[3]

To see how it deals with stylistic unacceptability, e.g. in the missing ear cases in (10), it suggests that unacceptability can result from (a) inadequate effects or (b) gratuitous effort. In other words, if a reformulation would cost the hearer less effort, then some element of stylistic infelicity should be perceived, but it will be forgiven if, e.g. the speaker had no time for reformulation.

[3]For further discussion of coherence-based approaches, see Moeschler & Reboul 1994; Moeschler 1997.

Relevance theory thus provides a relatively developed framework in which the results of detailed studies of specific aspects of comprehension can be integrated and explained. In the rest of this section, we will briefly sketch relevance-theoretic analyses of some of the bridging cases that were problematic for the truth-based and coherence-based approaches.

A unique feature of the relevance-theoretic comprehension process is that it involves a mutual adjustment of content, context and cognitive effects. The hearer, in interpreting an utterance, may have more or less specific expectations of cognitive effects (Sperber & Wilson 1998). An individual who has specific expectations of cognitive effects should pay attention to stimuli in his environment which are likely to yield these effects, whether or not these stimuli are intrinsically salient. Similarly, retrieval mechanisms may search for background assumptions that are not otherwise highly accessible, but which may allow the derivation of the expected effects. Thus, expectations of cognitive effect may alter the accessibility of candidate interpretations. Since the accessibility of interpretations affects their acceptability, it follows that an utterance judged unacceptable in isolation may be made more acceptable by manipulation of expected effects.

Consider from this perspective the problems raised by Erku & Gundel in connection with the examples in (11):

(11) a. ?A car went by. *The dog* was barking.
b. Don't go near that house. *The dog* will bite you.
c. Don't go near the car. *The dog* will bite you.

Erku & Gundel note that while (11a) is generally unacceptable in isolation, it would become more acceptable if understood as a denial of the prior utterance in (12):

(12) Dogs never bark in cars.

Relevance theory sheds some light on this example. Recall that one way of achieving cognitive effects is by contradicting and eliminating existing assumptions. When (11) is said in response to (12), this line of interpretation becomes highly salient. In order to achieve the expected cognitive effects, the hearer must use the assumption that there was a dog in the car, which in turn becomes salient as an implicated contextual assumption. Similar points apply to (11b) and (11c). In (11b), for example, the first part of the utterance raises the question of why the hearer should not go near the house. The second part of the utterance can be interpreted as an answer to this question on the assumption that the dog is near the house; this assumption in turn become salient as an implicated contextual assumption. More generally, the more specific the expected cognitive effects,

the more the interpretation process will be effect-driven, and the less the bridging assumptions will need to be independently accessible in advance of the interpretation process, e.g. as part of a frame. Conversely, the more accessible the intended bridging assumptions, the more the interpretation process will be effort-driven, and the less the hearer will need to have specific expectations about intended cognitive effects.

Discourse connectives provide another way of manipulating expectations about intended cognitive effects. The presence of a discourse connective can dramatically alter the interpretation of bridging examples, as the questionnaire results show. Thus, consider (34)-(37), where a discourse connective is added in the second member of the pair:

(34) a. I prefer Edinburgh to London. I hate the snowy winters. [London 100%]
b. I prefer Edinburgh to London. However, I hate the snowy winters. [Edinburgh 100%]

(35) a. I prefer Italy to England. I hate the pasta there. [England 80%]
b. I prefer Italy to England. However, I hate the pasta there. [Italy 100%]

(36) a. I prefer the town to the country. The traffic really bothers me. [country 80%]
b. I prefer the town to the country. However, the traffic really bothers me. [town 100%]

(37) a. Kevin moved from New Zealand to England. He loves the sheep. [England 100%]
b. Kevin moved from New Zealand to England. However, he loves the sheep. [New Zealand 80%]

In the relevance-based framework, discourse connectives are seen as guiding the search for relevance by constraining the choice of contexts and cognitive effects. Connectives like ‘but’ and 'however' encourage a 'denial of expectation' understanding (see e.g. Blakemore 1987; Moeschler 1989; Luscher 1989), in which the second part of the utterance is seen as contradicting and eliminating a potential implication of the first part. The addition of such a discourse connective thus encourages specific expectations of cognitive effects, in turn increasing the accessibility of any background assumptions needed to obtain these effects. In this way, the relevance-theoretic framework integrates the effort and effect factors, and shows how they interact.

According to Erku & Gundel, a relevance-based approach should shed some light on the stylistic unacceptability of examples like (7) above, in which a highly accessible interpretation had to be rejected on grounds of factual implausibility. In fact, the

questionnaire results suggest that subjects are not particularly worried by these examples, and do not usually mark them as unacceptable. Relevance theory suggests that an utterance should be found unacceptable if the speaker manifestly could not have intended it to satisfy the addressee's expectation of relevance. This may be for reasons of effort, as in (10c-d) above, where the obvious reformulations in (10a-b) would manifestly have spared the addressee some gratuitous processing effort. Or it may be for reasons of effect, as in (34a) above, where although 100% of subjects chose *London* as antecedent, 60% of them found the utterance unacceptable: presumably because it is hard to see how the speaker could reasonably have expected to convince them that she preferred Edinburgh to London on the ground that *London* had snowy winters. There should also be interesting interactions between effort and effect. For instance, consider (38), with the acceptability ratings shown in brackets:

(38) a. Peter went to a Japanese restaurant. The waitress was from Osaka. [5.77]
b. Peter went to a Japanese restaurant. The cashier was very friendly. [5.69]
c. Peter went to a Japanese restaurant. The poison was in the fish. [4.23]
d. Peter went to a Japanese restaurant. The ambulance came in 10 minutes. [2.46]

In the questionnaire, (38d) received a very low acceptability rating, presumably because the bridging assumption was considered too inaccessible given normal expectations of effect. Relevance theory predicts that this example could be made more acceptable either by increasing the accessibility of the bridging assumption, or by manipulating the audience's expectation of effect - for example, by producing the utterance as a one-liner in a Woody Allen stand-up routine. Results of this type would confirm the relevance-theoretic linkage between effort and effect.

Having sketched the assumptions of relevance theory and their application to some of the questionnaire examples, we would like to end by making some more direct comparisons between relevance-based and coherence-based approaches.

4 Comparing coherence theory and relevance theory

Let us start with an example that we have used several times, as presenting problems for both truth-based and coherence-based approaches:

(18) I ran from the classroom to the playground. The children were making too much noise.

This utterance has two interpretations, on both of which the second part of the utterance can be seen as giving an Explanation for the event described in the first part. This presents a problem for the coherence-based approach, since knowing the intended coherence relation is not sufficient for reference resolution.

A paper on disambiguation by Asher & Lascarides (1993) suggests a possible solution. Where an evidential relation is involved, the hearer should choose the interpretation which provides most evidence for the expected conclusion. Consider their example (39):

(39) a. They put a plant there. It ruined the view. [prediction: 'plant' = factory]
b. They put a plant there. It improved the view. [prediction: 'plant' = tree]

The word *plant* is ambiguous. In this case, knowledge of the intended coherence relation is not sufficient for disambiguation, because whichever interpretation is chosen, the second part of the utterance will describe a Result of the event described in the first part. A hearer in this situation, according to Asher & Lascarides, should choose the interpretation that 'reinforces coherence relations': that is, the one most likely to lead to the result described, the one that provides most evidence for an expected conclusion. On this approach, according to Asher & Lascarides, *plant* in (39a) should be taken to mean 'factory' because a factory is most likely to ruin the view, while in (39b) it should be taken to mean ‘tree’ because a tree is most likely to improve the view. By the same token, *children* in (18) should be taken to mean ‘the children in the playground’ or ‘the children in the classroom’, depending on which interpretation is seen as providing most evidence for the claim that the speaker ran from the classroom to the playground.

This approach is very similar to the strong, global truth-based approach discussed at the end of section 2.2.1, and shares many of the problems listed there. It also raises empirical problems for both coherence-based and truth-based approaches. In the first place, it suggests that the hearer of (39) will be unable to disambiguate the word *plant* correctly until the end of the following utterance. All the psycholinguistic evidence shows that tentative disambiguations are made on-line, but neither the coherence-based approach nor the strong global truth-based approach sheds any light on how this process works. By contrast, the relevance-theoretic comprehension procedure described in (30) above is compatible with the psycholinguistic evidence that hearers follow a path of least effort in the course of on-line disambiguation. It also sheds some light on how expectations of specific cognitive effects may alter the accessibility of interpretations.

A second empirical problem with Asher & Lascarides’ proposal is that hearers do not always choose the interpretation that ‘reinforces coherence relations’, i.e. that provides the most evidence for an expected conclusion. Consider (40):

(40) a. Many young people nowadays are unemployable. They can't even write a letter.
b. John wrote a letter. His wife was unimpressed.

The phrase *write a letter* can mean either 'write a letter of the alphabet' or ‘engage in correspondence'. In (40a) and (40b), it would be understood to mean ‘engage in correspondence’. However, this is not what Asher & Lascarides’ disambiguation strategy predicts. Each of (40a) and (40b) has two interpretations which fit the same evidential coherence relation: Explanation in (40a) and Result in (40b). According to Asher & Lascarides, the hearer should choose the interpretation that provides the most evidence for the expected conclusion. In (40a), he should therefore ask himself which would provide more evidence for the fact that young people are unemployable: that they can’t write a letter of the alphabet, or that they can’t engage in correspondence? Similarly, in (40b) he should ask himself which would be more likely to leave John’s wife unimpressed: the fact that he wrote a letter of the alphabet, or the fact that he engaged in correspondence? In both cases, Asher & Lascarides’ strategy would pick out the ‘letter of the alphabet’ interpretation, and in both cases, this prediction would be wrong.

Relevance theory explains why. According to the relevance-theoretic comprehension procedure, hearers should follow a path of least effort in computing cognitive effects, and stop when their expectation of relevance is satisfied. For a variety of reasons having to do with both effort and effect, the most accessible interpretation of the phrase *write a letter* is 'engage in correspondence'. In the first place, it would normally be taken for granted that someone can write a letter of the alphabet; as a result, it is hard to see how this interpretation could be intended to achieve adequate effects in normal circumstances. By contrast, it is easy to see how the information that someone had engaged in correspondence might be intended to achieve adequate effects, and hearers are likely to develop a ready-made encyclopaedic schema giving easy access to these standard effects. Frequent use of the phrase in this interpretation should further increase its accessibility. Thus, the relevance-theoretic comprehension procedure gives some insight into the interpretation of these examples.

A further problem with coherence-based approaches to pragmatics is that they provide no account of how isolated utterances are understood. Consider (41), for example:

(41) John wrote a letter.

According to Asher & Lascarides, this example would be disambiguated using ‘domain

information'; when (41) is embedded in a longer discourse, this domain-based interpretation could be overridden by considerations of coherence. The reference to 'domain information' would presumably involve some version of the truth-based pragmatic criterion, and would be subject to the same criticisms. Notice, also, that the two senses of *write a letter* stand in an entailment relation (one cannot engage in correspondence without writing a letter of the alphabet) so that the interpretation on which (41) means that John wrote a letter of the alphabet is always the more likely to be true. A strong truth-based criterion would therefore run into particular problems with this example. By contrast, the relevance-theoretic comprehension strategy works as well for (41) as it does for the longer discourses in (40).

Notice, too, that the coherence-based approach would have to provide some account of how to draw the line between truth-based and coherence-based processes. This is by no means straightforward, as is shown by the examples in (42):

(42) a. John's a very very clever man.
b. John's a very clever man, very.
c. John's a very clever man, a very clever man.
d. John's a very clever man. He's a very clever man.

All these utterances involve repetition: in (42a) and (42b) of a word, in (42c) of a phrase, and in (42d) of a sentence. It would be nice to be able to offer a unitary account of how repetitions achieve their effects, but this is not possible on a coherence-based approach. On this approach, repetition of full sentences, as in (42d), is generally seen as involving a coherence relation of Elaboration, or Restatement. However, coherence relations only hold between units of discourse, and the question arises of how many units of discourse there are in (42a-c). It seems implausible to claim that mere repetition of a word, as in (42a), would introduce a separate unit of discourse. Hence, coherence-based approaches must provide separate accounts of repetition in single units and pairs of units, and must furthermore provide a cut-off point between (42a) and (42d).[4]

By contrast, the relevance-theoretic approach provides a straightforward account of all cases of repetition. Repetition involves additional processing effort. On relevance-theoretic assumptions, this extra effort must be offset by extra effects, and according to the

[4]In fact, coherence theory imposes a second arbitrary cut-off point between repetitions of full sentences which are adjacent in discourse and those separated by other utterances. Generally, coherence relations are only seen as holding between adjacent units of discourse. Coherence-based approaches must thus provide separate accounts, not only of intra-sentence repetitions, but also of intra-discourse repetitions that are not adjacent.

relevance-theoretic comprehension procedure, the hearer should follow a path of least effort in computing these extra effects. Thus, the hearer of (42a-d) would probably assume that John is cleverer than he would otherwise have thought; (and since each successive utterance in the series demands greater effort, it would be interesting to investigate whether the level of expected effect increases accordingly). In this case, as in previous ones, the relevance-theoretic framework seems to explain a wider range of examples, more accurately and with less machinery, than the coherence-based approach.

5 Conclusion

This paper had two main aims: first, to shed some light on the analysis of bridging, and second, to compare three recent approaches to pragmatic analysis, based on notions of truth, coherence and relevance. We have tried to show that truth-based approaches are inadequate on their own, and that coherence-based approaches must generally be supplemented by some form of truth-based approach. Coherence-based approaches thus suffer from all the problems of truth-based approaches, and more besides. We have tried to show that the relevance-theoretic approach is superior to the coherence-based approach on both descriptive and theoretical grounds: it applies as well to isolated utterances as to longer stretches of discourse; it provides a comprehension procedure which does not involve computing all possible interpretations and choosing the best one; it shows how expectations of relevance can affect the accessibility of candidate interpretations, and how accessibility can in turn affect acceptability; and it explains most of the intuitions of acceptability that coherence-based approaches were designed to explain. We would like to end by commenting briefly on the relations between the truth-based and relevance-based approaches.

Relevance theory, unlike most other approaches to pragmatics, has no independent maxim or convention of truthfulness. It does, however, assume that the assumptions entertained by an individual may vary in their strength (i.e. degree of evidence), and that the strength of input assumptions affects the strength of any cognitive effects. The stronger the cognitive effects, the greater the relevance. It follows that expectations of relevance may lead to expectations about the strength of input assumptions, i.e. abut the factual plausibility of the proposition expressed by the utterance, and of the contextual assumptions needed to derive the expected cognitive effects. Since we have seen that expectations of relevance may affect the accessibility of candidate interpretations, it also follows that considerations of factual plausibility may affect the accessibility of bridging assumptions and candidate reference resolutions. Thus, the relevance-theoretic framework

should be able to explain the intuitions about factual plausibility that underpin the truth-based approach.[5]

References

Asher, N. & Wada, H. 1988 A computational account of syntactic, semantic and pragmatic principles for anaphora resolution. *Journal of semantics* 6: 309-344.

Asher, N. & Lascarides, A. 1993 Lexical disambiguation in a discourse context. *Journal of semantics* 12: 69-108.

Blakemore, D. 1987 *Semantic constraints on relevance*. Blackwell, Oxford.

Clark, H. & Haviland, S 1977 Comprehension and the Given-New Contract. In R. Freedle (ed.) *Discourse production and comprehension*. Ablex, Norwood, NJ: 1-40.

Erku, F. & Gundel, J. 1987 The pragmatics of indirect anaphors. In J. Verschueren & M. Bertuccelli-Papi (eds) *The pragmatic perspective*. J. Benjamins, Amsterdam: 533-45.

Giora, R. 1997 Discourse coherence and theory of relevance: Stumbling blocks in search of a unified theory. *Journal of pragmatics* 27: 17-34.

Grosz, B. 1981 Focusing and description in natural language dialogues. In Joshi, A., B. Webber & I. Sag (eds) *Elements of discourse understanding*. Cambridge UP.

Grosz, B. & Sidner, C. 1986 Attention, intentions and the structure of discourse. *Computational linguistics* 12: 175-204.

Hawkins, J. 1978 *Definiteness and indefiniteness: A study in reference and grammaticality prediction.* Croom Helm, London.

Hobbs, J. 1979 Coherence and coreference. *Cognitive science* 3: 67-90.

Hovy, E. 1990 Parsimonious and profligate approaches to the question of discourse structure relations. In *Proceedings of the 5th international workshop on natural language generation*: 128-136.

Lewis, D. 1983 Scorekeeping in a language game. In D. Lewis (ed.) *Philosophical papers*, vol. 1. Oxford UP, Oxford.

Luscher, J. 1994 Les marques de connexion: Des guides pour l'interpretation. In J. Moeschler, A. Reboul, J. Luscher & J. Jayez (eds) *Langage et pertinence: Reference temporelle, anaphore, connecteurs et metaphore*. Presses Universitaires de Nancy, Nancy: 175-227.

Mann, W. & S. Thompson 1988 Rhetorical structure theory. Toward a functional theory of text organisation. *Text* 8: 243-8.

Matsui, T. 1995 *Bridging and relevance*. University of London Ph.D. thesis.

Moeschler, J. 1989 *Modélisation du dialogue: Représentation de l'inférence argumentative.* Hermès, Paris.

Moeschler, J. 1997 Les relations entre événements et l'interprétation des énoncés. In J. Moeschler et al. (eds) *Pragmatique de la référence temporelle*. University of Geneva research document 3, vol.

[5]For further discussion of the relations between truth and relevance, see Sperber & Wilson (forthcoming).

1: 203-30.
Moeschler, J. & Reboul, A. 1994 *Dictionnaire encyclopédique de pragmatique.* Paris, Seuil.
Reboul, A. 1994 L'anaphore pronominale: le problème de l'attribution des référents. In J. Moeschler, A. Reboul, J. Luscher & J. Jayez (eds) *Langage et pertinence: Référence temporelle, anaphore, connecteurs et métaphore.* Presses Universitaires de Nancy, Nancy: 105-173.
Reboul, A. et al. 1997 *Le projet CERVICAL: Representations mentales, référence aux objets et aux événements.* CRIN-CNRS and INRIA-Lorraine.
Reinhart, T. 1981 Pragmatics and linguistics: An analysis of sentence topics. *Philosophica* 27: 53-94.
Samet, J. & Schank, R. 1984 Coherence and connectivity. *Linguistics and philosophy* 7: 11-35.
Sanders, T. Spooren, W. & Noordman, L. 1992 Towards a taxonomy of coherence relations. *Discourse processes* 15: 11-35.
Sanders, T., Spooren, W. & Noordman, L. 1993 Coherence relations in a cognitive theory of discourse representation. *Cognitive linguistics* 4: 93-133.
Sanford, A. & Garrod, S. 1981 *Understanding written language*. John Wiley, Chichester.
Sanford, A., T. Moar, K. & Garrod, S. 1988 Proper names as controllers of discourse focus. *Language and speech* 31: 43-56.
Sidner, C. 1983 Focusing in the comprehension of definite anaphora. In M. Brady & R. Berwick (eds) *Computational models of discourse*. MIT Press, Cambridge, MA: 267-330.
Sperber, D. & Wilson, D. 1995 *Relevance: Communication and cognition* (2nd edition). Blackwell, Oxford.
Sperber, D. & Wilson, D. 1998 The mapping between the mental and the public lexicon. In P. Carruthers & J. Boucher (eds) *Thought and language.* Cambridge University Press, Cambridge: 184-200.
Sperber, D. & Wilson, D. forthcoming Truthfulness and relevance in communication. Paper delivered to the San Marino conference on Grice's heritage, 1998.
Wilson, D. 1998 Discourse, coherence and relevance: A reply to Rachel Giora. *Journal of pragmatics* 29: 57-74.