QUANTITY MAXIMS AND GENERALISED IMPLICATURE[1]

Robyn Carston

Abstract

A theory of generalised conversational implicatures (GCIs) is being developed by Stephen Levinson and others working within the Gricean tradition in pragmatics. In this paper I begin to examine the components of this theory: the maxims and heuristics proposed, the schema for resolving conflicting predictions of these heuristics and the logical rules for generating GCIs. I compare this account with an alternative pragmatic theory, one which follows from the Relevance approach to communication developed by Deirdre Wilson and Dan Sperber. This differs from Gricean treatments in many respects, of which the following are particularly significant in the present context: (a) it employs a single pragmatic principle; (b) it makes no categorial distinction between generalised and particularised implicatures; (c) it does not employ a system of default inference rules to generate implicatures; (d) it makes a principled distinction between explicit content (what is said) and implicatures. I find this latter approach more adequate, in all of these respects, in accounting for the examples of pragmatic inference standardly called generalised implicatures. My broader claim is that a Relevance-based pragmatics stands a better chance of providing a psychologically sound explanation of utterance interpretation, since it is grounded in a theory of how we process and represent information in general, and is not construed as a module of special rules, based on an (unmotivated) analogy with generative grammar.[2]

1 Conflicting quantity maxims

The starting point for a major strand of work in pragmatics over the past decade is two examples of generalised conversational implicature which Grice gives in his seminal paper 'Logic and Conversation' (the examples are slightly adapted here for the sake of brevity):

---

[1] I am very grateful to Billy Clark and Vladimir Żegarac for their help and encouragement. I would also like to thank the students in my 1990 'Issues in Pragmatics' class, especially Janis Beavon and Agnieszka Urbanowicz, for their lively discusssion of various issues raised in this paper.

[2] Many of the matters discussed deserve more consideration than they get here. This is very much a paper which presents work in progress and which needs to be extended. In particular I do not consider in any detail the important work of Larry Horn on generalised conversational implicature, an omission soon to be rectified.

(1) a. I saw a beauuiful house yesterday.
b. The house was not my own.

(2) a. I broke a finger.
b. The broken finger is my own.

(1a) would 'normally (in the ABSENCE of special circumstances)' carry the implicature given in (1b) and (2a) would standardly carry the implicature in (2b). Both of course are cancellable without giving rise to a semantic contradiction, a central characteristic of all aspects of utterance meaning which are pragmatically derived rather than linguistically decoded. Within the Gricean system (1b) is claimed to follow from the first maxim of quantity, given in (3), which sets a lower bound on the quantity of information encoded by the linguistic form employed. (2b), on the other hand, is derived via the second maxim of quantity, also given in (3), which sets an upper bound on (semantic/linguistic) informativeness.

(3) **First maxim of quantity**: Make your contribution as informative as is required (for the current purposes of the exchange).
Second maxim of quantity: Do not make your contribution more informative than is required.

Now both Horn (1984, 1988) and Levinson (1987, 1988) have claimed that these two kinds of application of Gricean principles are in 'fundamental conflict'. To see this consider the entailment relations in (4).

(4) a. 'This is my house' entails 'This is a house'
b. 'This is my finger' entails 'This is a finger'

In each case the sentence on the left is 'semantically stronger' or more informative than the one on the right which it entails. The clash then is as follows: In (2) we have an inference from a semantically weaker proposition (regarding an unspecified finger) to a semantically stronger one (narrowing possible reference of 'a finger' to one of the speaker's own). In (1) we have an inference from a semantically weaker proposition (regarding an unspecified house) to the negation of the semantically stronger one (regarding speaker ownership). The point can perhaps be made most clearly by considering what the implicature of (2) would be if it worked the same way as the implicature in (1), that is, if it was derived from the first maxim of quantity. The reasoning based on this maxim would involve the assumption that since the linguistic content of the utterance is as informative as necessary, any more informative utterance (that is, one carrying more entailments) is not intended, and its propositional content can be assumed not to hold. The implicature would then be: 'The broken finger is not my own', that is, it would be a contradiction of the implicature derived when inferencing is based on the second maxim of quantity, the interpretation which answers to our intuitions about the example.

There have been many reformulations of these two quantity maxims, for instance by Atlas & Levinson (1981), Horn (1984, 1988) and Levinson (1987, 1988), in an attempt to see what is really going on here. I'll concentrate here on the recent work by Levinson, leaving an examination of Horn's version to later work. The two principles in (5) are taken from Levinson (1987, 67-68):

(5) **Q-principle**

a. *Speaker's maxim*: don't provide a statement that is informationally weaker than your knowledge of the world allows.

b. *Recipient's corollary*: take it that the speaker made the strongest statement consistent with what he knows.

I-principle

a. *Speaker's maxim: the maxim of minimisation*
'Say as little as necessary' i.e.produce the minimal linguistic clues sufficient to achieve your communicational ends.

b. *Recipient's corollary: enrichment rule*
'Amplify the informational content of the speaker's utterance, by finding a more specific interpretation, up to what you judge to be the speaker's m-intended point.'

In Levinson (1988) these are referred to as inferential heuristics which have developed as 'ways of over-coming the constraints of the narrow bandwidth of human speech' (1988, 4), rather than as maxims which follow from a principle of speaker/hearer cooperativeness. Levinson's idea seems to be that the behavioural maxims are derived from the inferential heuristics hearers use in interpreting utterances. The inferential heuristics corresponding to the two principles are given in (6):

(6) a. **Q-heuristic**: What isn't said, isn't the case.
b. **I-heuristic**: What isn't said is the obvious.

Familiar examples of the two kinds of implicatures claimed to be inferred on the basis of these principles/maxims/heuristics are given in (7) and (8):

(7) **Q-implicatures** (the negation of a semantically stronger proposition than the one expressed by what is said):

a. Mary failed some of her exams. <all, some>
implicates: Mary didn't fail all of her exams.

b. We scored three goals. <...four, three, two, one>
implicates: We scored no more than three goals.

c. John is a poet or a philosopher. <and, or>
implicates: John isn't both a poet and a philosopher.
- K (John is a poet)
- K (John is a philosopher)
(where ' - K (p) ' is to be understood as 'the speaker doesn't know whether or not p')

d. Her dress was red. {red, blue, green ...}
implicates: Her dress wasn't red and blue, etc.

(8) **I-implicatures** (semantically stronger than what is said):

a. Susan turned the key and the engine started.
   implicates: The engine started as a result of Susan's turning the key.
b. If you mow the lawn I'll give you £10.
   implicates: If and only if you mow the lawn I'll give you £10.
c. Paul's book is good.
   implicates: The book Paul read/wrote/borrowed/bought, etc, is good.
d. She was locked in a cell. The window was barred.
   implicates: The cell had a window.
e. John walked into the room and he laughed.
   implicates: John walked into the room and he, John, laughed.

The examples in (7) are the classical scalar types, depending on the assumed ready availability (high salience) to the hearer of (some version of) the 'Horn' scales shown in angle brackets beside examples a, b and c, or some other set of contrasting terms, as in example d. The second and third implicatures in (7c) also depend on these scales of items ordered with respect to semantic strength but are usually called clausal implicatures. They arise for the utterance of a sentence involving use of an item (here 'or') which does not entail an embedded sentence (for example, 'John is a poet') when use of a higher item in the scale (here 'and') would entail the embedded sentence (see Gazdar 1979, Levinson 1983). I'll concentrate on the scalar cases here. Given the assumed salience of the appropriate scale and given the Q-principle, the reasoning assumed to be involved in deriving the implicature is as follows: 'since the speaker has chosen a weaker, less informative expression where a stronger one of equal brevity is available, he would be in violation of the Q-principle if he believed the stronger statement to hold. Therefore, he must believe that the stronger statement does not hold, and he has done nothing to stop me thinking that he so thinks, therefore he implicates the negation of the stronger statement.'[3]

The I-implicatures in (8), on the other hand, are a rather heterogeneous set of cases. The crucial property they have in common is that 'they are more specific interpretations - what is implicated is a subcase of what is said' (Levinson, 1987, 66), and they are in potential conflict with Q-implicatures. Consider the 'pseudo Horn scale', <iff, if>, where the items are ordered in

[3] As this statement indicates the scalar implicatures should be generally construed as epistemically qualified, an issue not addressed here. Secondly, some of the scales could be more finely calibrated, for example, <all, most, many, some, few> in Levinson (1983, 134). This raises further epistemic considerations, since it's usually felt that the speaker endorses some of the possible scalar implicatures more strongly than others. Relevance theory easily accomodates these facts since it is built into its account of ostensive communication that communicated assumptions come with varying degrees of strength. So implicatures can be stronger or weaker, on a continuum of degrees of strength shading off into those assumptions a hearer derives entirely on his own initiative.

terms of semantic strength just as in the previous Q-implicature cases. If the Q-principle were applied here, using the sort of reasoning given just above, the implicature would be 'It's not the case that if and only if you mow the lawn I'll give you #10'. This is clearly not the right prediction, but what precludes it in this system of principles as so far outlined? What ensures that the I-principle operates here rather that the Q-principle? Various attempts at answering this question have been made, by Atlas & Levinson (1981), Horn (1984) and Levinson (1987), for instance. Before looking at one of the proposed solutions, though, we should briefly consider another set of cases which interact with the Q and I cases. These are the M-implicatures, which follow from the M-heuristic (which subsumes Grice's manner maxims concerning brevity and avoidance of obscure expressions). Levinson (1988)'s heuristic is given in (9), followed by some examples of M-implicatures:

(9) **M-heuristic**: what's said in an abnormal (or marked) way isn't normal

a. Mary caused the car to stop.
   implicates: Mary stopped the car in a nonstandard way, e.g.by using the hand-brake.
b. John employed an amanuensis to do his typing.
   implicates: The person John employed is male.
c. I have a female sibling.
   implicates: I don't have normal sisterly relations with her.

The idea here is that use of the phrase 'cause to stop' rather than the shorter, standard 'stop', the use of 'amanuensis' rather than the common high-frequency word 'secretary', and the use of the phrase 'female sibling' instead of 'sister', are marked (mannered, one might say) and so indicate something marked or non-standard about the action, entity, relationship, etc, that they represent. Horn and Levinson point out that there is an evident complementarity (or clash) between the I-principle, inducing enrichment of information in line with stereotypical expectations, and this M-principle, which blocks inference to the stereotypical, inducing instead implicatures regarding the unusual nature of what is being described.

Let's move now to the problem of how these three principles or heuristics interact, in particular to how the clash between the Q- and I- heuristics is resolved. I'll consider here just the 'schema' given in Levinson (1987), since it is briefer and neater than its forerunners and Levinson intends it as an amalgamation of the earlier ones.

(10) **Revised Resolution Schema** (Levinson, 1987, 71):

a. Genuine Q-implicatures from tight Horn scales and similar contrast sets of equally brief, equally lexicalized linguistic expressions 'about' the same semantic relations, take precedence over I-implicatures.
b. In all other cases, the I-principle induces stereotypical specific interpretations, unless:
c. there are two (or more) available expressions coextensive in meaning, one of which is unmarked in form and the other marked in form. In that case, the unmarked form carries the I-implicatures as per usual, but the marked form M-implicates the non-applicability of the pertinent I-implicatures.

What is going on here is an attempt to ensure that the Q-principle is given precedence, though with certain constraints on what constitutes a scale in order to block its application in a range of cases where it would give wrong predictions. So, for instance, there is no scale <iff, if>, since 'iff' is not 'equally lexicalised'. Nor are there scales, <since, and>, <regret, know>, leading in the latter case from 'I know that John is unwell' to 'I don't regret that John is unwell'; these are claimed to be ruled out by the ''about' the same semantic relations' requirement. Then the M-principle comes into play in the case of marked expressions, such as 'amanuensis' relative to 'secretary', 'caused x to stop' relative to 'stopped x', preventing, for instance, the standard I-implicature from 'secretary' to 'female'. Otherwise the I-principle has free rein, enriching informational content in accordance with uncontroversial, stereotypical assumptions. However, the 'potential explosion of I-implicatures' is further constrained by an undefined R(elevance)-principle (Levinson, 1987, 71) which presumably provides the hearer with a heuristic for determining the 'speaker's m-intended point' referred to above in (5) under the I-principle (b).

The heuristics and the resolution schema, for controlling which applies when, are central to Levinson's goal of formulating a theory of generalised conversational implicature (GCI), where this is taken to be a distinct phenomenon from particularised implicature. GCIs are default or preferred interpretations which are inferred quite generally unless there are specific contextual assumptions to defeat them. So an account of the logic involved in their derivation is needed to complete the picture. Levinson currently favours the adoption of some system of non-monotonic reasoning, probably the kind known as default logic. I consider this option below in section 4. But before that, I'd like to take a closer look, in the next two sections, at the I-principle/heuristic and at the alleged clash between it and the Q-principle/heuristic.

2 Q and I enrichments

A head-on clash between two members of a system of four principles can hardly be considered a desirable state of affairs. One theoretical impulse might be to seek a deeper principle underlying these which accounts for what, given this deeper principle, is just a superficial appearance of tension. Faced with the data of (7) and (8) this would be the route I would try since, as Reinhart (quoted in Levinson, 1987, 122) points out, both the Q- and the I- principles 'induce inferences more informative than what was said'. The unwieldy nature of the resolution schema, given in (10), should induce greater efforts in this direction. As it stands it is little more than a list of the problems and an imposition of an ordering on the heuristics: Q and M have precedence over I and the application of Q is strictly circumscribed. At best, it's observationally adequate[4] and no-one with an interest in psychological reality could be happy

[4] It's not clear that even this level of adequacy is achieved. Given the constraints on the sorts of items that can form genuine scales, it's difficult to think of examples for which a Q/I clash remains a possibility so that the precedence of the Q-principle over the I-principle can come into force. Example (2a) 'I broke a finger' might be such a case, but, if it is, the resolution schema seems in fact to

to stop here. There is no explanation of why the ordering should go this way and there is no mention of whether and, if so, how, this schema is actually mobilised by a hearer in the process of utterance interpretation. It's clearly a hangover from Gazdar (1979)'s formal solution to the projection problem for 'potential presuppositions and implicatures'. As such it's an approach based on a methodology developed for generative grammar, taking as its data speaker/hearers' judgments about properties of sentences rather than utterances and entirely ignoring the psychological facts of on-line interpretation. It is a subtheme of this paper, taken up more explicitly at the end, that this direct importation of the methodology of grammar into pragmatics is neither necessary (for lack of any other) nor desirable.[5]

Rather then tinker about with the resolution schema I'd like to go back and look more closely at the grounds for proposing the schema, that is the very existence of the clash itself. Let's take a closer look at the I-implicatures given in (8) above. Consider the format of example (8b), the 'if' to 'iff' example, and compare it with the format of (7a), both given above exactly as in Levinson (1987). The 'some' case in (7a) has essentially the format in (11):

(11) what is said: at least some of the ...
what is implicated: at most some / not all of the ...
what is communicated: just some / some but not all of the ...

The numeral case in (7b) works similarly with what is said being 'at least three', what is implicated 'at most three' and the overall significance (what is communicated) 'exactly three'. So also the disjunction in (7c), where the inclusive disjunction is what is said and the implicated assumption allows for at most one of the disjuncts, so that what is communicated is that one and only one of the disjuncts holds, that is, it's an exclusive disjunction. Compare this now with (8b) and consider why this isn't set up as in (12):

(12) what is said: if P then Q
what is implicated: if not P then not Q
what is communicated: iff P then Q

This is surely the parallel format with (11). Alternatively, the presentation of (7a) could be changed to parallel that of (8b), in a move directly from 'at least some' to 'just some'/'some but not all'. Now, what are the implications of this seemingly trivial uniformising of presentation? The point is that the appearance of a clash evaporates when this is done. In both sorts of case what the hearer understands as communicated is exactly what the schematic forms in (11) and (12) indicate as 'what is communicated'. In both cases there is an entailment

---

make the wrong prediction. On the assumption that there is some scale or salient set of contrasts to account for the Q-implicature in (1a), the same scale, whatever it is, should apply just as well to (2a) and the Q-implicature, 'The finger was not my own' should take precedence over the opposing I-implicature, 'The finger was my own'. The picture so far sketched looks observationally inadequate with regard to the implicatures of indefinite descriptions.

[5] See Dinsmore (1982) who puts these points clearly and convincingly in his review of Gazdar (1979).

relation between this and the proposition indicated as 'what is said', so in both cases there is a move from a semantically weaker or less informative proposition to a semantically stronger or more informative one. Furthermore, in neither case is the proposition given as 'what is implicated' stronger or weaker than the proposition given as 'what is said', that is, there is no entailment relation between them. So whichever format is ultimately favoured, that in (12) or that of (8b), the outcome is that either the implicature and what is taken to have been said by the utterance do not differ in relative strength or informativeness in either case, or the implicature is 'what is communicated' and it is more informative than what is said (it entails what is said) in both types of case. So in fact the scalar quantity cases and at least two of the I-implicatures cases, conjunction buttressing and conditional perfection, as they are known, work in essentially the same way. I will consider the other I examples below, since they seem to be a bit different.

The hearer's enrichment rule, in (5), taken by Levinson to follow from the I-principle, seems then just as adequate in accounting for the Q cases as for the I cases, since all of these involve the hearer in a process of building on the decoded content (what is said, in this sort of framework) to find a more specific interpretation. If this is right the separate Q-principle is redundant. There is no reason to suppose that because the natural language formulation of these richer interpretations involves a negative in some cases but not in others that this is grounds for setting up two distinct rules. However, while this enrichment rule is a move in the right direction its operation has to be constrained: '...I-implicatures can be as rich and specific as seems consistent with (a) an unreduced maxim of Relevance, and (b) what, given the assumed state of mutual knowledge, the speaker might reasonably have m-intended.' (Levinson, 1987, 72). Furthermore, the I-principle is not just this enrichment rule; it has been formulated in a range of ways and from a number of perspectives. In addition to those given in (5) and (6) Levinson (1987, 66) mentions the following from Atlas and Levinson (1981): 'the 'best' interpretation of an utterance is the most informative one consistent with what is non-controversial', and the convention of non-controversiality includes an assumption 'that referents and situations have stereotypical properties'. It's clear in Levinson (1988) that (undefined) notions of 'obviousness' and 'stereotypicality' are central to the notion of enrichment he envisages.

In the next section I'll try to show (a) that the remaining examples in (8) do not work in the way he assumes and that there is reason to doubt that some of them are implicatures at all, and (b) that for at least some of the I examples there is nothing 'obvious' or 'stereotypical' about the enrichments induced. I will then suggest that given a properly defined Relevance principle, the I-principle and any further mechanism for determining m-intentions is unnecessary. To the extent that 'obviousness' and 'stereotypicality' do play a role in the explanation of these examples, they follow from a properly developed notion of relevance.

3 Informativeness and relevance

Let's take a look at the remaining 'I-implicature' cases in (8). The genitive construction in (8c) encodes the fact that there is some relation between Paul and the book and perhaps constrains that relation to some extent, but particular

contextual assumptions are essential in order to fix the exact nature of the relation, as the implicature given by Levinson makes quite clear with its list of alternatives: 'read/wrote/borrowed/bought, etc'. Levinson wants to develop a theory of generalised implicatures which accounts for them as preferred or default interpretations but there is no reason to suppose that any one of the listed relations is the default interpretation either for the genitive construction in general or for the particular instantiation of it here, 'Paul's book'. It's no more obvious or stereotypical that the book is one Paul read than that it's one he's just bought or one that he gave to Mary, and so on. The idea of default interpretations is generally problematic (see sections 4 and 5 below) but for this example it's not just problematic, it's a complete non-starter. What seems to be involved here is not stereo-typicality or any broader notion of the non-controversial but some concept of relevance or contextual appropriateness; that is, the I-principle doesn't apply in any of its manifestations. Given the context-dependence of the inference here it looks as if it should be reclassified as a particularised implicature rather than a generalised one. However, if the implicature is 'The book Paul wrote is good', what is the explicit content of this utterance (what is said) ? It must be something like 'The book, in relation x with Paul, is good', where x is a variable. But this is semantically incomplete (that is, not truth evaluable) and so, arguably, not what Grice intended by his category of 'what is said'. Furthermore, on this analysis, it's not ultimately part of the interpretation of the utterance, simply serving the function of giving the hearer access to the alleged implicature, which subsumes it. All of this runs counter to the strong intuition that such an utterance does have explicit content, that it's not a case where what is communicated is just an implicature and that the contextual fixing of the relation between Paul and the book is simply part of what's involved in arriving at a 'full identification of what the speaker has said' (Grice, 1975, 44).

However, Levinson is quite consistent in his view that any assumption involving pragmatic inference is an implicature, as is evident from example (8e), where he takes the assumption derived by assigning the referent of 'John' to 'he' to be an implicature. Surely an assignment of reference to 'he' is necessary in the process of determining what has been said by an utterance of (8e) (its explicit content). Given that there are no grammatical restrictions on the reference of this pronoun and no particular contextual assumptions containing individual concepts which might function as sources of referents, the most accessible, indeed the only available, referent for 'he' is whatever the referent of 'John' is. This is simply the best we can do given the lack of any further clues. What this has to do with implicature is far from obvious; certainly it's a pragmatic matter, as is disambiguation and the recovery of ellipsed material, both of which may be necessary to arrive at what is said. Just where the saying/implicating distinction lies and the extent of pragmatic enrichment required at the level of what is said are taken up briefly in section 7.

(8d) is one of various cases of 'bridging implicature' Levinson includes in the class of I-implicatures. Now the given implicature 'The cell had a window' is certainly not stronger, in any sense, than either of the propositions expressed at the level of explicit content in this example, supposedly the defining characteristic of I-implicatures. However, let's assume, by analogy with the pronominal reference case in (8e), that what Levinson really means here is that what is implicated is 'The window of the cell was barred'. Then the points already made with regard to (8e) arise again: 'The window was barred' is

arguably not a complete propositional form since, among other things, it requires assignment of reference to the definite description, before we arrive at what the speaker has said. A more coherent position on (8d) is to take the implicature to be the bridging sort as originally given, that is, it's an implicature which is accessed in order to provide a referent for 'the window', an essential part of arriving at what is said. On this view of things the implicature is not a subcase of what is said, the supposedly standard property of I-implicatures, and so does not clash in any way with the predictions of the Q-principle. This sort of bridging phenomenon is clearly context-dependent, here depending on the mention of a cell in the first part of the utterance and the general knowledge which that gives the hearer access to. Therefore, on the Gricean characterisation of implicature, it is particularised rather than generalised.

The conclusion that emerges from this examination of the examples in (8) is that there is no general class of implicatures generated by the I-principle which have the properties Levinson ascribes to them and which potentially clash with the output of the Q-principle. In fact, in at least some of these cases we don't seem to be dealing with implicature at all.

As we've seen, there are various (non-equivalent) formulations of the I-principle, involving notions such as obviousness and stereotypicality. 'What isn't said is the obvious' cannot distinguish I-enrichments from Q-enrichments, at least not without some explanation of the concept of obviousness. Many of the scalar cases would seem to fall in here: it's standardly obvious when a speaker uses 'some of the x' that he is not committing himself to 'all of the x'. As for the idea that I-enrichments tend to be the 'best interpretation to the stereotypical, most explanatory reading' (Levinson, 1988, 5), this just seems to be false for at least some of the standard cases of I-implicatures:

(13) a. He opened the door and she handed him the key.
b. If you don't pick up the garbage I'll make you a cake.

(13a) implicates that she handed him the key after he opened the door; this implicature of temporal ordering is one of the standard conjunction buttressing cases. But it is surely not an enrichment of content along stereotypical lines; if anything, it is the exact opposite of the typical scenario which forms part of a readily accessible mental script about opening doors. Nevertheless, the inference is readily drawn and the stereotypical ordering (first the key, then the opening of the door) is inaccessible. The I-principle cannot explain this case, but this is just the sort of example it was set up to deal with. The same point applies to (13b): in the absence of contrary assumptions, the implicature 'if you do pick up the garbage I won't make you a cake' is inferred. Again this runs strongly counter to intuitions about what is stereotypical or obvious in human transactions. This implicature might lead to the accessing of further assumptions, such as that the speaker must think that her making a cake is undesirable to the hearer. But the point is that this 'fallacy of denying the antecedent' type of implicature seems to go through whether it fits with stereotypical expectations or not. I don't want to deny that there is a role for some notion of stereotypicality in accounting for the assumptions we construct and the inferences we draw in processing new information, but whatever role it plays it doesn't offer an explanation for these cases of I-implicature. The inescapable conclusion here is that neither 'obviousness' nor 'stereotypicality' do the job of accounting for the sort of enrichments exemplified by the so-called I-implicature

cases and nor do they clearly distinguish the Q and I cases.[6]

The existence of an I-principle in addition to a principle of relevance (which Levinson grants is necessary) is highly questionable. Grice, himself, said of the second maxim of quantity, which is the forebear of the I-principle: 'its effect may well be secured by a later maxim, which concerns relevance' (Grice, 1975, 46). He, of course, didn't develop any such maxim or principle though he recognised its centrality to the sort of system he was setting up. In line with this Horn (1984) takes 'the clash' to be between a Q-principle (much as Levinson's), accounting for the scalar cases, and what he calls an R-principle: 'Say no more than your hearer needs', which appears to be understood as doing the work of both Levinson's I- and R- principles. Now this bears a clear relation to that Principle of Relevance which is the backbone of the pragmatic theory formulated by Sperber and Wilson (SW). In the most straightforward cases the guarantee of optimal relevance, which all utterances communicate according to this principle, is a guarantee of an adequate range of contextual effects and of minimal processing effort required to derive these effects. Saying no more than your hearer needs follows from the relevance requirement of not putting him to any unnecessary processing effort. Horn's principle raises the question of what constitutes a hearer's needs; say no more than your hearer needs for what? The answer to this follows from the SW principle of relevance: 'say no more than your hearer needs in order to achieve the intended set of effects'.[7]

Similarly, I believe, the SW principle of relevance encompasses the territory taken to belong to the I-principle. This could be demonstrated in two ways: by showing that the concepts of the I-principle such as enrichment to the obvious or stereotypical, when they work, follow from optimal relevance, and, secondly, by giving relevance analyses of the implicature data which has motivated the setting up of these principles. I make a start on the latter task in section 6. As to the former, I won't explore it here, except to point out the plausibility of the view that stereotypical assumptions, such as that 'birds fly', 'secretaries are female', 'you stop a car by using the foot brake', etc, are made readily accessible by the very words used in the respective utterances: 'bird', 'secretary', 'stopped the car'. Given the minimal processing effort side of the guarantee of optimal relevance, the frequency with which such assumptions enter into interpretations of utterances which encode these concepts is predicted. Of course contextual factors affect accessibility: a discussion about penguins, for example, would tend to demote the accessibility of the assumption about birds flying. Relevance theory assumes that encyclopaedic information is differentially accessible at any given point in interpretation, and that the accessibility ranking

---

[6] Obviousness is clearly a context-dependent notion, and as such one that doesn't seem appropriate for an account of default inferences. Stereotypicality might seem more promising in this respect, but this is not so sure either, given the work of Barsalou, summarised in Barsalou (1987), showing the context dependent nature of stereotypicality.

[7] Harnish (1976) extends the maxim system still further, having, for instance, a maxim enjoining representationalness: 'in so far as possible, make your sayings 'mirror' the world', with various sub-maxims of this. I believe it won't be difficult to show (though it has yet to be done) that these fall out from the SW principle of relevance.

is constantly changing as context evolves. However, Levinson favours an account of the frequent high accessibility of certain assumptions, stereotypical ones, in terms of default logical rules, an idea which I turn to now.

4 Default logical rules

There are three components to the 'theory of generalised conversation-al implicature' which Levinson envisages: the inferential principles/ heuristics, the resolution schema, and a system of default logical rules. So far he has been inexplicit about the latter but he has stressed two properties that generalised implicatures (GCIs) have and which must be reflected in the logic involved in deriving them: (a) they are default inferences, that is, they arise unless something specifically blocks them; (b) they are defeasible or cancellable (a property shared by all implicatures, whether generalised or particularised). It follows from this latter property that the logic employed must be non-monotonic, that is, it must allow for the possibility of additions to a set of premises leading to the dropping of conclusions generated by the unaugmented set of premises.A comparative example should make this clear:

(A) Premises:	All humans are fallible
	Pat is a human
Conclusion:	Pat is fallible
(B) Premise:	Pat has eaten three of the cakes
Conclusion:	Pat hasn't eaten four of the cakes

(A) is a typical example of a deductive inference; no matter what further premises are added to the given premise set (provided none of the existing ones is erased) the set of conclusions can only increase, that is, it's monotonic. In (B), a typical case of a GCI, on the other hand, the addition of a premise can lead to the dropping of the conclusion. Consider, for example, adding as a premise 'If Pat has eaten three of the cakes she's bound to have eaten more'. According to Levinson (1989, 466) the inference in (B) 'falls into the category of non-monotonic default inferences now being energetically explored in Cognitive Science' and he refers to the work of Raymond Reiter among others. Reiter (1980)'s default logic consists of rules with the following format:

(14) $\dfrac{\alpha(x): \quad \beta(x)}{\beta(x)}$

This should be read in the following way: if $\alpha(x)$ holds and if $\beta(x)$ is consistent with the data-base, then you can infer $\beta(x)$.

To take the much cited case of Tweety, the idea is that given the information that Tweety is a bird, we infer, in the absence of information to the contrary, that Tweety flies. The specific rule employed would look something like the following:

(15) BIRD(x): FLY(x)

FLY(x)

That is, 'if x is a bird and it can be consistently assumed to fly, then you can infer that x flies'. The inference is blocked if there is information in the database (or context) with which the assumption that Tweety flies is inconsistent, for example, that Tweety is a kiwi, or that Tweety's wings are broken.

Levinson's idea seems to be that the same sort of rules might be involved in deriving generalised conversational implicatures. Then the default rule in (16) would be responsible for the scalar implicature generally inferred from the use of 'some' and the rules in (17) for the temporal and causal implicatures associated with the use of 'and'.

(16) γ(some): γ(not all)

γ(not all)
(where γ is a simple sentence frame)

(17) a. P and Q: P before Q

P before Q

b. P and Q: P caused Q

P caused Q

With regard to the rules in (17) they will give the result we want for (8a) above. Assuming there is no specific information in the current context at odds with the assumptions that Mary turned the key before the engine started and that her turning the key caused the engine to start, these inferences go through. However, let's consider some further cases of conjunction:

(18) a. I did the dishes and gave the baby a bath.
b. It's 30 C in Delhi and 35 C in Bombay.

In processing an utterance of (18a) the default rule for temporal ordering would go through in accordance with the general under-standing that the baby bathing followed the doing of the dishes. However, we don't want the causal conclusion, which the second rule will give us unless it is blocked by information to the contrary, but what information prevents it going through here ? Certainly we can construct the uncontroversial assumption that an instance of doing dishes doesn't cause an instance of bathing a baby. But we would be accessing such an assumption solely to prevent the working of this rule; without such a rule the construction of this negative information wouldn't be necessary and, intuitively at least, it isn't necessary and it isn't constructed. The situation gets worse when we consider (18b) for which negative contextual assumptions must be constructed in order to block both of the rules which would lead to the conclusions that it's being 30 C in Delhi preceded and caused it's being 35 C in Bombay. Furthermore, there are other conjunction buttressing implicatures,

such as the teleological one mentioned by Levinson: 'Mary turned the key in order to start the engine', and others which capture different sorts of temporal and consequential connections that arise when we broaden the range of examples:

(19) a. She was baking a cake and listening to radio 4.
b. He slept deeply and dreamed that he was flying.
c. We investigated the problem and it was far more complex than we expected.

Whatever the temporal relations are between the events or processes described by the pairs of conjuncts in (19a) and (19b), it's clearly not the case that they are simple sequential cases. In (19a) it may be that the two processes coincided while in (19b) the dreaming span is most likely fully contained in the sleeping time. In (19c) there is some kind of consequential relation between the investigation and the complexity of the problem but it's not a causal one. Since these connections will 'generally, in the ABSENCE of information to the contrary' be inferred, it looks as if the advocate of default logic rules is driven here to setting up some further rules for 'and', in addition to the two given in (17). Within any such expanded set of rules for 'and' there will be rules which are at odds with each other, for example, the temporal sequence rule and the temporal containment rule; at this point, the very idea of a default interpretation for 'and' seems to collapse. And, returning to (18b), our case of symmetric 'and', understanding this would require a hearer to construct a plethora of negative contextual assumptions in order to prevent a range of quite mad applications of the set of default rules. These observations merely underline the general problem with this whole approach to implicature: it is inappropriately generative and rigid, that is, it ignores the context-dependence of all pragmatic inference.

With regard to the rules given in (17), Hirschberg (1985, 44) mildly remarks: '...under this characterisation one would be forced to say that the 'normal' use of 'and' is the asymmetric understanding ...[in cases where we don't want the asymmetric understanding] some unspecified 'special circumstances' contextually cancel this 'normal' temporal/causal understanding. But why asymmetric 'and' should be more 'normal' than symmetric 'and' is unclear. And how these canceling circumstances may be identified is never discussed by those who assume their existence.' Sensibly, she rejects any derivation of implicatures by a default logic, asserting that there is no principled basis upon which to assign defaults.

A couple of other considerations mediate against this approach, though further counter-evidence is hardly necessary. Recall that the conjunction buttressing implicatures depend on the I-heuristic: stereo-typically, an engine starts when the (appropriate) key is turned, etc. Levinson clearly believes that this heuristic must be employed in conjunction with a logic and so isn't sufficient by itself to account for the implicatures derived. The question arises though of whether any such heuristic is necessary once a full-blown default logic is set up. Given such rules as those in (17), rules which operate whenever the particular lexical item occurs unless they are specifically blocked on the particular occasion, what role is there for the I-heuristic ? It's clear that this heuristic is completely redundant; any force it might have had has been taken over by the apparently inevitable and automatic operation of the default logical

rules.

It's a well-aired fact that all the various informational enrich-ments that conjunctions can support also arise for non-conjoined juxtapositions of sentences. So the examples in (20) would all tend to implicate the same sorts of connections between the events or processes they describe as do their conjoined counterparts:

(20) a. Mary turned the key. The engine started.
b. He slept deeply. He dreamed he was flying.
c. We investigated the problem. It was far more complex than we expected.

There is no lexical item to hang the default rule on in these cases and without setting up default rules for every possible pairing of sentences in the language (obviously an impossible task) this approach to generalised conversational implicatures seems to have nothing to say about these cases although they are essentially identical to the 'and' cases.

Many of the above points apply also to the scalar cases, resting on the Q-heuristic, and there are further problems with these, most of which I will not pursue here. Briefly though, let's consider the implicatures dependent on the cardinal number scale. The standard neo-Gricean position on these is that the semantic (decoded) content of a cardinal number term, 'n', is taken as its lower bound, equivalent to 'at least n', and the preferred/default interpretation of 'exactly n' is derived by combining this semantics (what is said) with the scalar implicature which gives the upper bound 'no more than n'. If the default logic idea were carried over to this case there might be a rule of the sort given in (21):

(21) α(n): α(no more than n)

α(no more than n)
(where 'n' is a cardinal number term)

This is simply at odds with the facts: a cardinal 'n' may be understood in at least the following three ways: 'exactly n', 'at most n', and 'at least n', and which of these is taken to have been communicated depends on context. Consider the example in (22), adapted from Hirschberg (1985, 92). In the context of A's question, C's response is most likely to be understood as 'at least £5', assuming that movies don't usually cost more than £5. However, as a response to B's question, C would most likely be understood to communicate 'at most £5'.

(22) A: Can you afford the movies?
B: Do you have £10 ?
C: I have £5.

It's difficult to see what contextual information here prevents the application of the rule in (21), whose conclusion, in interaction with the assumed semantic content of the numeral would generate the 'exactly' understanding. Surely the point is that one interpretation rather than the others is more appropriate or relevant in the context and there just is no default interpretation. This treatment

of the cardinals is seriously flawed in the further respect that there is actually no way to derive the 'at most £5' reading from the semantics of numerals here assumed, whether default rules are used or not. To arrive at this one-sided upper bounding interpretation the lower bounding semantics which constitutes what is said (explicit content) has to be overridden, but there are no grounds for doing this, no maxim flouting or contextual inconsistencies.

Finally, a treatment of generalised implicatures in terms of default rules cannot be extended to encompass most of the M-implicature cases, since (as in (9a) and (9c)) their derivation frequently depends on the use of a phrasal circumlocution when a lexical item with roughly the same extension is available. Since the phrasal possibilities are infinite there is no point in trying to set up rules for particular cases. Nor can more general meta-rules covering classes of cases be concocted; any attempt along these lines would ultimately be forced to such a level of generality that it would amount to a restatement of the M-heuristic itself.

Many of these problems for the default logic approach arise precisely because what the rules do (and what any such rule system cannot avoid doing) is associate generalised implicatures with particular lexical items when what they really depend on in many instances is whole sentences or propositions. Levinson is by no means alone in associating implicatures with particular lexical items or with classes of lexical items (as in the scalar cases). Horn (1972), Gazdar (1979) and Harnish (1976) show the same tendency when they discuss 'generalised' cases though they are less explicit. Hirschberg (1985, 43-44) sees Grice's original characterisation of generalised implicature as the source of this mistake. She says: 'The relative context-independence of generalised conversational implicatures has led to serious confusion over the distinction between this phenomenon and the conventional force of an utterance. Even Grice (1975, 56) describes generalised conversational implicatures as cases in which 'the use of a certain form of words in an utterance would normally (in the ABSENCE of special circumstances) carry such-and-such an implicature or type of implicature'.' Certainly there is something misleading about this emphasis on 'a certain form of words' but a bit of close attention to Grice's text leads me to doubt that he intended that these implicatures would arise in this general fashion whenever particular lexical items or classes of lexical items were used, irrespective of their sentential context. His way of introducing examples was standardly as follows: 'Anyone who uses a sentence of the form 'X is meeting a woman this evening' would normally implicate ...' and when discussing particularised cases he says 'there is no room for the idea that an implicature of this sort is NORMALLY carried by saying that p' where p is something propositional (truth evaluable). The implication in this is that he takes generalised cases to depend on the entire linguistic content of the utterance, that is, its logical or propositional form. If this is right he would certainly not have favoured the use of default logical rules which hinge on lexical items or classes of lexical items for deriving these implicatures.

5 The generalised/particularised distinction

The failure of this attempt to treat one class of conversational implicatures as default inferences casts serious doubt on the distinction between generalised and particularised implicatures. This distinction is treated as fundamental by

the neo-Griceans (Atlas, Gazdar, Horn, Levinson) and Levinson (1989) explicitly favours two distinct theories of implicatures, one for the generalised cases and one for the particularised cases. In her excellent study of scalar implicature Hirschberg (1985, 42) rejects the distinction and calls it 'an artifact of the inventiveness of analysts - or lack thereof'. In the context of her discussion of asymmetric 'and' she says of the temporal, etc, cases discussed above: 'These generalised implicatures, while more context-independent than particularised implicatures ...are still context-dependent' (43). And in her later discussion of the sorts of examples identified by Levinson, Gazdar and Horn as 'scalar quantity implicatures' she shows that they do not form a natural class and are subsumed by her broader category of scalar implicature which includes implicatures with varying degrees of dependence on the specifics of particular contexts.[8]

It's interesting to note that if we were to take the generalised/ particularised distinction as absolute and were entirely literal in applying Grice's characterisation of the former, many cases of metaphorical and hyperbolic utterances, as he describes them, would fall into the class of generalised implicatures. Consider his example 'You are the cream in my coffee' which has the 'metaphor interpretant 'You are my pride and joy'' (Grice, 1975, 53). He supplies no particular contextual assumptions against which to interpret the example and of course none is necessary. Since, as he says, the example involves a categorial falsity its literal meaning will be rejected as the intended interpretation and the putative implicature derived quite generally, across contexts. This will be the case for a wide range of relatively common uncreative metaphors, such as 'She's an angel', 'This room is a pigsty', etc.which communicate a standard implicature. Such implicatures which arise in this general cross-contextual way, in the absence of assumptions which block them, answer to the description of generalised implicatures. Nevertheless it's perfectly clear that Grice saw these as just as much cases of particularised implicature as his other examples of alleged maxim flouting, such as the irony case and the notorious testimonial case, whose implicatures depend on particular contexts which he explicitly sets up. This might look like an incoherence in his system; I would prefer to see it as a indication that the distinction is not to be taken too seriously. The force of the distinction is further dimmed when we think about the 'generalised' cases that fall under Levinson's M-heuristic (or Horn's Q-principle). Despite their fairly general occurrence across contexts these would have been particularised implicatures arising from the flouting of one of the manner maxims ('avoid obscurity of expression', 'be brief') in Grice's original scheme.

Sperber & Wilson (1987, 748) express doubt that the distinction was of theoretical significance for Grice and they themselves favour a continuum of cases of implicature, with some resting on very widely held and standardly available assumptions about the world, some resting on more specific cultural assumptions which are shared by a wide range of people through to these which

---

[8] Hirschberg (1985, 56) also points out the need for a notion of 'relevance in context' (1985, 56) to account for the salience of particular scales and of other non-scalar orderings in certain contexts. Again, it would be my aim to show that once this notion of relevance is properly developed it will account for the full set of examples she covers, without the need for an additional quantity maxim.

are dependent on very specific and transient information. On their relevance theory view, within which I am working, the same principles and logic will apply to all these cases. Levinson is right in his observation that work on implicature within relevance theory up to now has focussed on the sorts of cases Grice would have characterised as particularised. He may also be right in his contention (1989, 465) that Grice's central interest was in the examples he dubbed 'generalised' because of their role in adjudicating between different views on the semantics of such natural language items as connectives and quantifiers.[9] However he moves from these points to the strong (and, in my view, unsupported) assertion that relevance theory cannot account for the sorts of cases Grice would have characterised as generalised: 'There can be no way to get the universal regularities of GCI predictions out of a theory of 'nonce-inference'; relevance theory was simply not concocted to deal with GCIs, and should relinquish ambitions in that direction.' (Levinson, 1989, 466). In the next section I will make a beginning to the task of showing that the concepts of relevance theory do provide for an account of these cases of implicature and that, therefore, no theory of generalised conversational implicature distinct from a theory of particularised implicature is called for.

6 A relevance approach

I don't propose to give yet another outline of relevance theory here, but there are some central points to bear in mind. A basic claim of the theory is that human cognitive activity is driven by the goal of maximising relevance: that is, deploying its information processing resources in such a way as to derive as great a range of contextual effects (contextual implications, strengthenings, and eliminations) as possible for the least expenditure of effort. This is what motivates us at any given moment to pay attention to some phenomena rather than to others. We are generally willing to attend to utterances because an utterance comes with the presumption that it will afford the hearer an adequate range of contextual effects, that is, a presumption that it's worth the hearer's while to turn his attention to it. He will of course have to put some effort into processing it: decoding its linguistic content, accessing a context of assumptions to interact with it, and computing its effects on that context. The other half of the presumption of optimal relevance says that he won't have to expend any gratuitous effort in achieving the intended effects, that is, the effort required will be adequately rewarded. The speaker's corollary of this is that she should try to formulate her utterance in such a way that the interpretation which is both the most accessible one to the hearer and which gives a satisfactory yield of effects is the one she intended.

An explanation of the 'generalised implicature' cases is going to depend quite heavily on the least effort side of the principle of relevance. Let's first consider an utterance of a conjunction such as (23a) which is usually taken to

[9] Grice (1981, 185) says, of generalised implicatures: 'These are the ones that seem to me to be more controversial and at the same time more valuable for philosophical purposes, ...this notion of generalized conversational implicature might be used to deal with a variety of problems, particularly in philosophical logic, but also in other areas.'

implicate (I-implicate, in Levinson's system) temporal sequence, and causal and teleological connections, as in (23b), (23c) and (23d):

(23) a. Susan turned the key and the car started.
b. Susan turned the key and then the car started.
c. Susan turned the key and as a result the car started.
d. Susan turned the key in order to make the car start.

A hearer of (23a) is given immediate access to his general knowledge about turning keys and starting cars, part of a bunch of strongly held, frequently used assumptions about people's interactions with cars. The cost of accessing these sorts of assumptions is negligible and they represent the kinds of relations which are quite standardly relevant to us (have contextual effects) since they enable us to make sense of events in the world, in particular human behaviour, and to predict that behaviour and its consequences. There is not a lot more to be said about this example, artificially abstracted from context as it is: the assumption that the car's starting is an outcome of Susan's action is simply the easiest assumption to access which offers the possibility of some effects to offset the effort of processing that piece of information. In the setting of a real conversational exchange, which would make accessible a range of assumptions in addition to these scripted ones, the utterance would be bound to have more effects.

So far then we're pretty much in line with Levinson's 'stereo-typical' enrichments version of the I-principle. But let's return to example (13a), repeated here as (24), an example which, as we saw in section 3, the I-principle fails to explain:

(24) He opened the door and she handed him the key.

The point of this example is that the temporal sequencing which is bound to be assumed is one of the opening of the door preceding the handing over of the key. The question then is why (24) isn't enriched to 'he opened the door after she handed him the key' in accordance with our much used door-opening script. (I don't want to imply that this is completely impossible but it's a very much less accessible interpretation, to put it mildly.) It begins to look as if there are two different ways by which we reach temporal sequencing assumptions: (1) via mental scripts, (2) via order of presentation of information. The two frequently work in tandem and make the same temporal ordering prediction, as in (23a). Order of presentation apparently holds sway in (24) and there is no sense of tension since scenarios in which a person opens a door in order to hand something to another person are not particularly unusual, although less acessible perhaps than the key-opening-the-door routine, all other things being equal, which they're not in this case. A serious conflict between these two routes accounts for the weirdness of the well-known examples in (25a) and (25b):

(25) a. He rode into the sunset and jumped onto his horse.
b. He went to bed and took off his boots.

These sorts of examples motivated Grice's original manner maxim enjoining orderliness and Harnish's 'mirroring' maxims6. However, as H.Clark (cited by

Gazdar 1979, 44) first noted, sequencing maxims simply make the wrong predictions about examples such as (26a) and (26b), where the temporal and consequence relations can go either way; these differ from examples (26c) and (26d) for which the 'backwards' relations are not possible:[10]

(26) a. He hit her. She screamed at him.
b. She screamed at him. He hit her.
c. He hit her and she screamed at him.
d. She screamed at him and he hit her.

The situation we find ourselves in is one where existing Gricean analyses simply can't cover the data: Levinson's I-principle fails for any temporal enrichment that isn't in line with stereotypical assumptions, such as (24); sequencing maxims make the wrong predictions for the interpretation of some of the non-conjoined cases.

Given some very intuitive observations about ease of processing effort, relevance theory can explain the full range of cases. It's just a fact about our minds that we find it easiest to process information when it is sequentially ordered. Imagine that the instructions for assembling a model aeroplane or for knitting a jumper were presented to you in random order or that the frames of a comic strip were not sequentially ordered. You would have considerable difficulty in understanding and, if you ultimately succeeded, the cognitive gains would be incommensurate with the effort expended. There is something deeply obvious about this; perhaps the explanation lies in some very basic facts about simple perceptual experiences: if an event x in the world immediately precedes an event y and if one sees these events, one generally sees event x before one sees event y and processes the latter perception in the context of the former. What is clear is that this is a cognitive matter and does not require the setting up of special pragmatic maxims; the facts about utterance interpretation follow from the facts about cognition. So the explanation for the ordering we assume for (24) lies with our bent to keep processing effort to a minimum. The experience of oddity we get in processing (25a) and (25b) arises because of the conflict between the ordering assumed on grounds of ease of processing and the script that the examples give access to. In these cases there is extreme difficulty in trying to construct a scenario which conforms to the ordering that least effort processing makes immediately accessible.

What about (26a) and (26b) then ? If this ease of processing explanation is right, how is it that we may understand (26a) as communicating that his hitting her was a consequence of her screaming and thus that the event described second preceded the event described first ? This interpretation is not possible for the corresponding conjunction in (26c). The explanation for this must lie with the unit being processed for relevance; in the non-conjoined cases the first sentence is processed individually and the bid to establish its relevance may raise questions regarding causes and results (among other possibilities): why did he hit her ? what was the consequence of his hitting her ? A readily

[10] This section owes much to Wilson (1990) who discusses the inadequacy of sequencing principles and of special-purpose principles in general. She makes a very strong case for a relevance treatment of a wide range of data concerning temporal interpretation.

available schematic mental script involving a man hitting a woman allows for a range of possible causes and consequences of this behaviour, the woman screaming being a candidate for both. The second sentence in (26a) may be the speaker's response to the question she judges her first sentence to have raised or it may simply reflect the way in which the information has come into focus in her own mind and so be, in effect, a slightly less well planned version of (26d). This option doesn't arise for the conjunction in (26c) which is presented as a unit of information so that the ordering assumptions which follow from ease of processing prevail.[11]

I hope this account is reasonable enough to counteract Levinson's assertion that Relevance theory has nothing to say about general tendencies in pragmatic inference. As was clear from the discussion in sections 2 and 3, the examples standardly given as I-implicatures do not form a natural class but work in quite distinct ways and there is no reason to expect identical explanations for them, though, in each case, I contend, the principle of relevance will play a crucial role. I won't attempt here an analysis of all the I-examples, but will briefly consider (8e), the preferred local coreference case. It looks like a very straightforward matter of assigning the most readily accessible referent to the pronoun. Certainly it isn't an enrichment that can be explained by appealing to stereotypical assumptions. It does seem 'obvious', in line with Levinson's characterisation of the I-heuristic, so perhaps it follows just as easily from both the principle of relevance and the I-heuristic. However, consideration of some other examples of pronoun reference assignment, of an equally 'generalised' nature, favours the Relevance account in terms of effort and effect over the appeal to an undefined notion of 'obviousness'. According to psycholinguistic work (cited in Garnham (1985, 151)) subjects standardly take the referent of 'he' to be John in (27a) and to be Bill in (27b) and they take longer to understand (27b) than (27a):

(27) a. John sold Bill his car because he hated it.
b. John sold Bill his car because he needed it.

The preferred reference assignment in (27a) conforms with the 'parallel function strategy', which is one of a set of well supported processing strategies given in the psycholinguistic literature: that is, the antecedent of a pronoun plays the same role in its clause (e.g.subject or object) as the pronoun does in its. The preferred reference assignment for (27b) does not conform with this strategy so some other principle must be involved in assigning reference here. Levinson's 'preferred local coreference' following from 'what isn't said is the obvious' cannot explain these examples. Let's go along with the parallel function strategy as a plausible strategy developed as an effort-saving device in interpretation. In accordance with this, the first interpretive hypothesis concerning pronominal reference for (27a) is the one which links the pronoun to 'John'. This is consistent with the principle of relevance in having an adequate range of effects; at least there is no reason to suppose that it won't, which is as far as we can go in considering these decontextualised examples. However, the corresponding first accessed hypothesis for (27b) (using the parallel function strategy) won't yield an adequate range of effects as it is

[11] Blakemore (1987, 117-123) discusses this sort of case in some detail.

contradicted by strong general assumptions about the reasons people have for selling things. So the other possible antecedent, 'Bill', is tried and is consistent with relevance expectations. An account along these lines not only explains the preferred interpretations but even offers an explanation for the greater difficulty subjects experience in understanding (27b) than (27a).

We also want a relevance account of the Q-implicatures, so-called. These have occupied a central place in neo-Gricean theory and are probably its strongest case. The following comments do not constitute a developed relevance analysis but are just some of the considerations that must inform such an account. Let's take the classical scalar implicature case, given in (7a) and repeated here:

(28) a. Mary failed some of her exams.
b. Some of the arguments are convincing.

An example of this sort is put forward in every book on pragmatics as the generalised conversational implicature par excellence. Yet the move from 'some' to 'not all' is actually quite restricted, as the following range of examples indicates:

(29) a. Some birds flew past the window.
b. I'd like some apples please.
c. It took me some time to get here.
d. The park is some distance from my house.

Firstly, there seems to be a distinction between two uses of 'some', the one pronounced /s m/ and the other pronounced /s m/, the former occurring in (29a) and (29b) where there is no pragmatic inference to 'not all'. However, even if we confine our attention to the /s m/ use of 'some', examples like (29c) and (29d) show that the occurrence of the implicature is still far more limited than a default rule would predict. It is open to the default advocate to claim that there is some contextual assumption blocking the inference in these cases, but no particular contextual assumptions are given and the point about generalised implicatures is that they should go through in the absence of such. Of course it's perfectly plain that there is something quite different about this use of 'some' from its use in (28), but, whatever the difference is, it is a difference which prevents the application of the default rule attached to 'some' and how it does this needs to be specified.

The sort of example usually given to illustrate scalar implicature is a simple atomic assertion with the partitive genitive construction, as in (28a) and (28b). These involve quantification over specific domains which are contextually determined, as required by the definite descriptions, 'her exams' and 'the arguments'. What is 'generalised' about this sort of example is the instruction, that the definite description encodes, to retrieve or set up a domain which accords with its descriptive content and to quantify in a certain way over it. Standardly such contexts are ones in which a contrast between 'only some' and 'all' is relevant, where a crucially different range of contextual effects will arise depending on whether the whole domain or merely a proper subset of the domain is involved. The 'at least some' (a subset which may or may not be a proper one) possibility is simply too vague in most contexts and so won't fulfil the first half of the principle of relevance, the guarantee of an adequate range

of effects. In fact the relevant contrasting quantitative concepts in such a case would standardly be the set consisting of not only 'all' and 'only some' but also 'none'; it will generally make a difference in terms of effects whether Mary failed all her exams, just some of them or none of them. Now these concepts don't form a Horn scale of the entailment sort that is standardly invoked in explaining such examples, but as Hirschberg (1985) has convincingly demonstrated nor do the vast range of cases that, for her, fall under the Q-implicature. What is at issue in each case is the relevant set of contrasts in the particular context and whether or not these are in an entailment relation with each other is beside the point.

This narrowing of the concept is not inevitable though, even when a domain of quantification is established; sometimes the broader concept is relevant enough:

(30) A: If any of the students have failed I'll be in trouble.
B: I'm afraid some of them have.

Here what matters to the hearer, A, is whether or not one or more students have failed; if so, it follows that A will be in trouble, from which, presumably, follows a rash of other implications. This is a threshold case and the upper bound is of no significance. I don't think there should be any inclination here to say that B's utterance implicates that not all of the students have failed. But there is no contextual assumption, say 'All the student have failed', which would block this supposedly default inference. There is no such inference because it's not required for the hearer to satisfy himself that the presumption of optimal relevance has been fulfilled.

I would argue that the enriching of the vague concept encoded by 'some' to 'some but not all', when it occurs, is just one of many concept narrowings affecting the level of explicit content (what is said). A different sort of narrowing would standardly occur in interpreting an utterance of (29c): a span of time is always involved in a person getting from one place to another, so (29c) always expresses an unremarkable truism but one that few speakers would intend to communicate. A hearer would standardly take a speaker to have explicitly communicated that the time taken was longer than he might have expected, that is, of such a length that certain implications follow, say, that this is the reason he's late, etc. Similar remarks could be made about (29d). In the next section I will argue that many of the 'informational enrichments' standardly treated as implicated assumptions in Gricean theory are really cases of concept narrowing at the level of explicit content.[12]

---

[12] I defer a relevance account of the indefinite article examples, (1) and (2), which opened this paper. Grice (1975, 57) makes some remarks about these examples which are very suggestive of a relevance treatment: 'The only difficult question is why it should, in certain cases, be presumed, independently of information about particular contexts of utterance, that specification of the closeness or remoteness of the connection between a particular person or object and a further person who is mentioned or indicated by the utterance should be likely to be of interest. The answer must lie in the following region: Transactions between a person and other persons or things closely connected with him are liable to be very different as regards their concomitants and results from the same sort of transactions involving only remotely connected persons and things.'

Finally, and briefly, it should be reasonably clear how a relevance analysis of the M-implicature cases will go. Processing 'male sibling' involves more effort than processing 'brother': 'male sibling' is longer and 'sibling' is a less frequently used word than 'brother', and so less accessible. The greater effort involved in processing the phrase must be offset by compensatory effects and some (rather indeterminate) implicature/s arise. This might seem very similar to the sort of account Levinson gives but there is an important difference. His account is explicitly transderivational, that is, it involves the hearer in some sort of comparison between the phrase the speaker used and a lexical item she could have used to pick out the same entity in the world. From his recognition that she chose to expend more effort he concludes that she must want to communicate something about the unstereotypical nature of the extension. This account attributes to the hearer an implausible amount of reasoning about the speaker's intentions and her use of language. No such comparison need be made on the relevance account as the accessibility differences are simply built into the organisation of memory. Hearers do not concern themselves with speakers' effort but look for effects to offset their own effort. Again, this sort of account follows naturally from facts about how we process information in general and the presumption of optimal relevance that accompanies linguistic stimuli.[13]

7 Implicature or explicature?

There's a good case to be made that these Q- and I- enrichments are not implicatures at all, but are pragmatically derived aspects of what is said (explicit content), where this is understood as the truth conditional content of the utterance. I've already made this point regarding the assignment of reference to the pronoun in (8e) and the fixing of the variable relation encoded by the genitive in (8c), where the meaning inferred is clearly a part of 'what is said' as Grice originally outlined it (Grice, 1975, 44). But it extends to the other cases as well. Consider the following examples:

(31) a. If the old king died of a heart attack and a republic was declared Sam will be happy, but if a republic was declared and the old king died of a heart attack Sam will be sad.
(adapted from Cohen 1971)

b. It's always the same at parties: either I get drunk and no-one will talk to me or no-one will talk to me and I get drunk.
(from Wilson 1990)

[13] I have said nothing here about the Relevance theory account of the inferential processes involved in implicature derivation. The only logical rules employed are deductive, that is, monotonic, but these are just one component of the non-demonstrative reasoning system of hypothesis formation and confirmation envisaged by Sperber and Wilson. Other crucial aspects of the picture are the degree of accessibility of assumptions, the varying strengths of assumptions and the procedures for resolving contradictions. These independently motivated features of the system together give rise to non-monotonic effects, such as implicatures, without calling for any special system of non-monotonic rules.

c. If each side in the soccer game got three goals, then the game was a draw. (from Levinson 1988, 29)

d. AIDS vaccination will cause some mortality: but losing some of the population is better than losing some and maybe all of it. (from Levinson 1988, 24)

e. If you've broken a finger you'll have to see the doctor.

Cohen (1971) presented examples like (31a) and (31b) as evidence against Grice's implicature treatment of the temporal and causal connotations of 'and'. The point is this: on Grice's account 'P and Q' and 'Q and P' are truth-conditionally equivalent, so reversing the order of the conjuncts in a given sentence should make no difference to truth-conditions. But this is just false: (31a) is not contradictory at the level of truth-conditional content and (31b) is not a redundant repetition, as the implicature analysis would predict; (31c) is clearly true which would not be the case if 'three' was understood here as 'at least three' at the level of explicit content, and there is nothing anomalous about (31d) as there would be if the first 'some' were taken as merely setting a lower bound. In short, pragmatic enrichment contributes to truth-conditional content, that is, to what is said (explicit content).

Grice wanted to give an implicature treatment to various aspects of meaning that arise in the use of indicative conditionals, so that he could maintain a material implication semantics for these. He met up with what he took to be a major problem: '...if the affirmation of 'if p, q' carries an implicature, its denial has to be interpreted as the denial of an implicature.This principle does not appear to be acceptable.' (from the William James lectures, 1967, reprinted in Grice, 1989, 83). In fact this phenomenon is far more widespread than he realised, as the examples in (31) indicate, and leaves no option but to allow that some pragmatically derived meaning is part of explicit content rather than an independent implicated assumption. And once it is acknowledged that pragmatic principles are involved in disambiguation, reference assignment, and recovery of ellipsed material, etc, processes which, on Grice's own characterisation, are central in determining what is said, it's evident that an instance of recovering what is said without pragmatic inferencing, if it ever occurred, would be exceptional.

As we've seen, Levinson takes any and all aspects of utterance meaning derived pragmatically to be implicatures. Let's look again at some of his I-implicature cases, repeated here, to see what this means:

(8) c. Paul's book is good.
implicates: The book Paul read/wrote/borrowed/bought/etc, is good.

d. She was locked in a cell. The window was barred.
implicates: The cell had a window.

e. John walked into the room and he laughed.
implicates: John walked into the room and he, John, laughed.

There are two ways of interpreting what is going on here. The most sensible construal of the bridging case in (8d) is that the implicated assumption supplies material to be used in developing the decoded content of the utterance into the

truth-conditional content of the utterance. It is not clear that this is the way Levinson (1987) sees it, given his characterisation of I-implicatures as semantically stronger than what is said. On the second interpretation, which Levinson (1987) lends itself to, as is evident from (8c) and (8e), it looks as if the putative implicatures entirely take over from what is said and there is nothing communicated at the level of explicit content, contrary both to intuition and to the spirit of the original Gricean programme. The first (happier) construal prevails in Levinson (1988), where the dominant theme is the role of pragmatics in determining what is said; then the propbsition, 'John laughed' in (8e), for example, is part of what is said by this utterance and the pragmatic assigning of 'he' as coreferential with 'John' is just one of many pragmatic inferences involved in determining what is said.

On the Relevance account pragmatically determined aspects of what is said (explicature) are distinguished from implicated assumptions. Of the assumptions communicated by an utterance, the explicit ones are those which are developed from a logical form encoded by the utterance. The distinction between explicatures and implicatures is captured in derivational terms: the processes involved in deriving an explicated assumption take the schematic, truth-conditionally incomplete, output of linguistic decoding and supply it with inferred material, as required to fix variables and to enrich content, in accordance with the criterion of consistency with the principle of relevance. It needs to be argued in much more detail but the evidence of examples such as those in (31) makes it look very likely that, in many instances at least, the examples of pragmatic inference that have been classified as Q- and I-implicatures are really pragmatically derived aspects of explicitly communicated assumptions.

Now Levinson (1988) is impressed by examples such as those given in (31), which he refers to as 'pragmatic intrusions on truth-conditional content' and he gives a veritable corpus of such examples in that paper. He continues to call these implicatures, but allows that these are implicatures which constitute a part of truth conditional content, hence what is said. At this point it might seem that there is just a terminological difference between Levinson and Relevance theory: both acknowledge pragmatic input at the level of explicit content. However, what Levinson has yet to supply is any coherent distinction between those 'implicatures' which contribute to what is said and those which function as independent assumptions communicated by the utterance. The relevance theory characterisation of explicit content in terms of developments of a decoded logical form, constrained by the principle of relevance, offers a coherent distinction.[14]

Furthermore, the failure to distinguish these two roles of pragmatic inference in utterance interpretation, together with some other basic assumptions of neo-Gricean theory, leads to the problem known as 'Grice's circle': 'Grice's account makes implicature dependent on a prior determination of 'the said'. The said in turn depends upon implicature: it depends on disambiguation, indexical

[14] Carston (1988) is a more extensive discussion of the implicature/ explicature dstinction. In that paper and in Recanati (1989) a number of criteria are set up and tested in an attempt to find a general answer to the question: when does a pragmatically determined aspect of utterance meaning qualify as a conversational implicature and when does it qualify as constitutive of what is said (an explicature)?

resolution, reference fixing, not to mention ellipsis-unpacking and generality-narrowing. But each of these processes, which are prerequisites to determining the proposition expressed, themselves depends crucially on implicatures. Thus what is said seems both to determine and to be determined by implicature' (Levinson, 1988, 17).

The Gricean circle is the outcome of the neo-Gricean conception of semantics, pragmatics and the semantic/pragmatic distinction. The ideas here are that (truth-conditional) semantics is a component of the grammar or language faculty and that pragmatics is another module (or set of modules) whose input is the output of the semantics module. It follows that the semantic/pragmatic distinction should be in terms of a truth-conditional/non-truth-conditional distinction. But the evidence of examples such as (31) shows that pragmatics contributes to truth-conditions, hence the semantics/pragmatics circle. In fact there is nothing viciously circular in having pragmatic input to the level of explicit content together with semantic input to the derivation of implicatures, provided we make a few distinctions, including the explicit/implicit distinction as given above, and provided we observe some of the psychological facts of on-line interpretation rather than thinking in terms of modules outputting whole propositionally shaped units of meaning.

So let's make another distinction. Levinson, in line with a long tradition, equates linguistic meaning with truth-conditional meaning: 'The 'said' can be taken to be truth-conditional content, the proposition expressed, the output of the process of semantic interpretation, the proper domain of a theory of linguistic meaning' (Levinson, 1988, 17) (my emphasis). But why should the referent that gets assigned to the pronoun 'he' on some occasion of utterance be considered a matter of linguistic meaning ? This is tantamount to saying that this pronoun is a billion-ways ambiguous. What we know when we know the meaning of the lexical item 'he', what is stored in the 'mental lexicon', is knowledge of the way it delimits the field of possible referents in any context. It is one of a huge range of natural language expressions, including other pronouns, demonstratives, tenses, gradable adjectives, quantifiers, which are clearly truth-conditionally incomplete.

How to characterise the semantic content of these expressions is an interesting question, recently opened up by work in relevance theory.[15] What this work shows is that some lexical items make an incomplete contribution to truth-conditional content and some others make no such contribution at all. What this in turn makes clear is that linguistic semantics cannot be equated with truth-conditional semantics. The output of linguistic decoding is seldom, if ever, fully truth-conditional, but requires considerable 'filling in' by pragmatic processes in order to become the sort of representation which can be given a truth-conditional semantics. In short, two levels of semantics, and in fact two kinds of semantics, have to be distinguished. The first concerns the sorts of concepts and constraints encoded by linguistic expressions, part of the theory of

[15] Blakemore (1987) develops an extensive account of lexical items whose semantics is universally agreed not to contribute to truth-conditional content. Wilson and Sperber (this volume) emphasise the range of roles that lexical items can play in utterance interpretation, distinguishing in particular those that make a fully truth-conditional contribution from those, such as pronouns, which merely constrain the truth-conditional content.

the speaker/hearer's linguistic knowledge, and therefore a psycho-logical matter. The second concerns the relation between fully propositional mental representations and the real world conditions they represent, and this is not part of a psychological theory ('truth conditions are not in the head' as it is sometimes put). Levinson dismisses the Relevance account of the first kind of semantics, calling it the 'semantic retreat' position and invoking Lewis's adage, 'semantics without truth-conditions is not semantics'.

When he writes of semantic interpretation Levinson appears to mean the formal apparatus of model-theoretical semantics: 'Thus we cannot obtain a truth-value (a maximal extension) without first doing the pragmatics.In short, if we can find cases of implicaturally determined reference, then we show that the whole semantic apparatus that recursively defines intensions and extensions is reliant on pragmatic input.' (Levinson, 1988, 39). But whatever the rules are for determining the truth-value of a proposition they are not psychologically represented. Doing this second kind of semantics, which deals in truth-values, individuals and mathematical functions, is a distinct enterprise from giving an account of the conceptual and procedural semantics which is part of our native speaker knowledge of language.

In fact Levinson ultimately embraces the necessary distinction (though without drawing out its implications) in his search for a way out of Grice's circle: 'I believe that a theory ...which provides a level of representation to which both pragmatic and semantic processes can contribute and which is the level at which semantic interpretation is defined, does in fact offer a [coherent] picture' (Levinson, 1988, 72). 'The slate thus represents the semantic and pragmatic content of accumulated utterances, and it is this representation as a whole that is assigned a model-theoretic interpretation' (Levinson, 1988, 22). He recommends Discourse Representation Theory and File Change Semantics which, as systems with this desirable property, promise an account of the interaction of semantic and pragmatic processes in arriving at what is said. I won't consider these here; the point though is that he is now advocating two quite distinct levels of semantics, the modular 'semantic processes' referred to no longer constituting the domain of truth-conditions (hence not semantic at all, if we call in Lewis), and the level which admits of model-theoretic interpretation. The earlier dismissal of Relevance theory for its level of semantically incomplete logical form must be retracted.

What still needs to come is an awareness of the cognitive psychological underpinnings of processes of utterance interpretation. A full appreciation of these will lead to the abandoning of a further factor contributing to the Gricean circle problem: the treatment of pragmatics as a module, or set of modules, of distinctive rules or principles, receiving input from and sending output to semantic modules. This is an extension of the modular structure of generative grammar to a theory of meaning; I briefly consider the appropriateness of this extension in my concluding remarks. Sperber & Wilson (1986) and Wilson & Sperber (1986) have made a strong case against this and for a pragmatics which emerges out of facts about our general information processing capacities. The arguments in this paper, against the effectiveness of the clashing I-, Q- and M-principles, and the mistaken resort to default logics, should cast further doubt on this approach.

8 Concluding remarks

I have taken issue both with the very idea of a theory of generalised conversational implicature and with details of Levinson's version of such a theory. The details criticised are: (a) the view that there is a clash between maxims (and that this is desirable), (b) that the I-principle accounts adequately for the implicatures usually attributed to it, and (c) that the derivation of generalised implicatures involves a system of default logical rules. As regards the idea that a distinct theory of GCIs should be developed, this presupposes that the distinction between generalised and particularised cases is clear and principled, but there is simply no evidence for this. There is instead a range of cases of implicature, some relying on assumptions that humans in general subscribe to, some on assumptions that are held by certain social and other groupings of people, and others which rely on the particular characteristics of the situation which the conversationalists find themselves in. Therefore we should be looking for pragmatic principles and logical abilities of a quite general sort which are called into play in the derivation of the full range of implicatures and perhaps in other types of inferential activity as well.

There is one last issue I'd like to touch on here. Levinson, Horn and others see it as a virtue of their systems that they have several principles/maxims 'interacting' with one other. Sadock (1986, 87), for instance, expresses misgivings with the Sperber-Wilson approach simply on the grounds that they run a pragmatic theory on a single principle: 'In their attempt to reduce the number of operative principles of pragmatic theory to one, S&W are running the risk of introducing complexity as an artifact of not recognising the complexity as arising from the interaction of quite distinct, unitary, orthogonal principles.' He gives no example of this 'artifactual complexity' but it's clear that he feels some general uneasiness about the singularity of the principle of relevance. He is certainly not alone in this. Levinson (1987, 76) explicitly opposes any reduction of the Gricean maxims and says: '...it is most doubtful that a single principle like SWR [the principle of relevance] can hope to cover the range of both the I- and the Q- inferences, their clash and their resolution' and, more forcefully, in Levinson (1989, 465) he talks of the 'impossibility of reducing countervailing principles to one mega-principle'. I hope that I have succeeded in at least casting doubt on the truth of this assertion in section 3.

There is a tendency among these neo-Griceans towards making analogies with the principles of generative grammar which work together to account for a wide array of grammatical facts. Horn (1988, 115) says: 'Pragmatics itself may be viewed as internally modular and interactionist, in the sense that the conceptually distinct subcomponents (suborientations) of pragmatic analysis may be simultaneously called upon within a single explanatory account of a given phenomenon, just as autonomous but interacting grammatical systems may interact to yield the simplest, most general, and most comprehensive treatment of some linguistic phenomenon (cf.the deconstruction of passive in Chomsky 1982).' This is a bold and intriguing claim, but, at least with regard to the sort of principles considered here, there doesn't seem to be any evidence for it. The intricate and richly productive interaction of a few general principles is an impressive feature of the Chomskian approach to grammar; so, for example, grammatical subsystems such as the categorial component, the transformational component, case theory and bounding theory do 'conspire together' to account

for the passive construction, and any one of these rules or principles, say constituent movement, plays a role in the generation of a wide range of superficially different structures such as passive, raising constructions, ergative, relative clauses and wh-questions, for example. However, I find nothing bearing the remotest resemblance to this in the sort of interaction of principles outlined by the neo-Griceans. Each principle is set up to account for a discrete set of cases: the Q-principle accounts for the Q-implicatures, the I-principle accounts for the I-implicatures, etc. They do not interact to give rise to a range of phenomena; in fact they conflict in their predictions and require extraneous constraints, of an entirely unexplanatory and unexplained sort, to resolve this conflict. I don't think generative grammarians would find a comparable clash of grammatical principles acceptable. The point then is that the pragmatic systems proposed simply do not consist of rules/maxims/ principles interacting in the favoured manner. There is the further question of whether this sort of interactionism is what we should be looking for in trying to build an explanatory account of the data of pragmatics. Should we be aiming for a system made up of tightknit autonomous modules interacting with each other in the way that grammatical principles do. To come to any conclusions about this will require careful consideration and comparison of the nature of a grammatical theory and of a pragmatic theory, and of the sorts of psychological abilities they are constructed to account for. The current state of play, though, is that none of the systems of pragmatic principles currently in the field even approximates the sort of interacting modular structure of the government and binding theory of grammar.

Finally, I wouldn't want to deny the possibility that there are bundles of heuristics/strategies/schemas involved in pragmatic inferencing. It's widely assumed that there are accessing strategies such as the already mentioned parallel function strategy, for example. What I would want to show is that these are geared towards and follow from the general cognitive goal of maximising relevance, that is, of deriving as many effects as possible while keeping processing effort to a minimum. It would be an interesting exercise to go through the psycholinguistic literature and look at the various comprehension strategies proposed to see if they can indeed be explained in these terms.

References

Atlas, J. & Levinson, S. (1981). It-clefts, informativeness, and logical form: radical pragmatics (revised standard version). In Cole, P. (ed.) *Radical pragmatics*. New York: Academic Press.1-61.

Barsalou, L. (1987). The instability of graded structure: implications for the nature of concepts. In Neisser, U. (ed.) *Concepts and conceptual development*. Cambridge: Cambridge University Press. 101-140.

Blakemore, D. (1987). *Semantic constraints on relevance*. Oxford: Basil Blackwell.

Carston, R. (1988). Implicature, explicature and truth-theoretic semantics. In Kempson, R. (ed.) *Mental representations: the interface between language and reality*. Cambridge: Cambridge University Press. 155-181.

Carston, R. (forthcoming). Pragmatics and the explicit/implicit distinction. University College London.

Cohen, J. (1971). The logical particles of natural language. In Bar-Hillel, Y. (ed.) *Pragmatics of natural language*. Dordrecht: Reidel.
Dinsmore, J. (1982). Review of Gerald Gazdar Pragmatics: implicature, presupposition, and logical form. *J. of Pragmatics* 6. 39-48.
Gazdar, G. (1979). *Pragmatics: implicature, presupposition, and logical form*. New York: Academic Press.
Garnham, A. (1985). *Psycholinguistics: central topics*. London: Methuen.
Grice, H.P. (1975). Logic and conversation. In Cole, P.and Morgan, J. (eds.) *Syntax and semantics 3: speech acts*. New York: Academic Press. 41-58.
Grice, H.P. (1981). Presupposition and conversational implicature. In Cole, P. (ed.) *Radical pragmatics*. New York: Academic Press. 183-198.
Grice, H.P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University press.
Hirschberg, J. (1985). A theory of scalar implicature. Doctoral dissertation. University of Pennsylvania.
Horn, L. (1984). Towards a new taxonomy for pragmatic inference: Q-based and R-based implicature. In Schiffrin, D. (ed.) *Meaning, form and use in context (GURT '84)*. Washington: Georgetown University Press. 11-42.
Horn, L. (1988). Pragmatic theory. In Newmeyer, F. (ed.) *Linguistics: the Cambridge survey*, vol.1. Cambridge: Cambridge University Press. 113-145.
Levinson, S. (1987). Minimization and conversational inference. In Verschueren, J. and Bertuccelli-Papi, M. (eds.) *The pragmatic perspective*. Amsterdam: John Benjamins. 61-129.
Levinson, S. (1988). Generalized conversational implicatures and the semantics/pragmatics interface. Cambridge ms.
Recanati, F. (1989). The pragmatics of what is said. *Mind and Language* 4. 295-329.
Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence* 13. 81-132.
Sadock, J. (1986). Remarks on the paper by Deirdre Wilson and Dan Sperber. *Parasession on pragmatics and grammatical theory, Chicago Linguistics Society* 22, 85-90.
Sperber, D. & Wilson, D. (1986). *Relevance: communication and cognition*. Oxford: Basil Blackwell.
Sperber, D. & Wilson, D. (1987). Précis of Relevance: communication and cognition. *Behavioral and Brain Sciences* 10, 697-754.
Wilson, D. (1990). Pragmatics and time. Talk given to the MIT conference on Time in Language, March 1990.
Wilson, D.& Sperber, D. (1986). Pragmatics and modularity. *Parasession on pragmatics and grammatical theory, Chicago Linguistics Society* 22, 67-84.
Wilson, D.& Sperber, D. (1990). Linguistic form and relevance. This volume.