



Embedded Implicatures

François Recanati

► **To cite this version:**

| François Recanati. Embedded Implicatures. Philosophical Perspectives, Wiley, 2003. ijn_00000372

HAL Id: ijn_00000372

https://jeannicod.ccsd.cnrs.fr/ijn_00000372

Submitted on 8 Sep 2003

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Embedded Implicatures

François Recanati
Institut Jean-Nicod

I. Introduction

Conversational implicatures are a species of pragmatic implication¹; they are implications of an act of 'saying'. The speaker's saying that p implies that q (given the presumption that he or she respects the norms of conversation — Grice's 'maxims' — or the overarching 'Cooperative Principle'). For example, the speaker's saying that some students came, together with the premiss that s/he is well-informed and tries to be as informative as possible, implies that not all students came. (If all students had come, the speaker ought to have said so.) Insofar as the speaker overtly intends the hearer to recover those pragmatic implications of the speech act, they are part of what the speaker means, though not part of what the sentence means. In this way we can account for certain aspects of utterance meaning within pragmatics, without burdening semantic theory.

Conversational implicatures thus understood have two important features. First, they result from an inference. Now 'inference' can be used in two ways : the broad and the narrow sense. In the strict, narrow sense, inferences satisfy what I call the *availability condition* : whoever makes an inference (in that narrow sense) is *aware* that the judgment he or she arrives at is inferentially based upon some previous judgment. No such condition applies to inferences in the broad sense. Imagine someone hearing the doorbell, and coming to believe, on that basis, that there is someone at the door. This example involves both an inference in the broad sense and an inference in the narrow sense. Identifying the sound one hears as that of the doorbell arguably involves an 'inference', as some cognitive scientists tell us; but the subject is not aware that that is so. The availability condition is not satisfied.

¹ By 'pragmatic implication', I mean the implication of an action.

On the other hand, even though the subject's inference that there is someone at the door is spontaneous and unreflective, still it is 'available' to the subject, who knows that the basis for his or her judgment is the fact that the doorbell is ringing. This is an inference in the narrow sense, one that takes place at the personal rather than merely at the sub-personal level.

Conversational implicatures, I take it, are inferences in the narrow sense: the subject knows that his or her judgment relative to what the speaker implies is based upon some independent judgment regarding what the speaker says. To take a standard example: if, when asked whether I can cook, I reply 'I am French', my utterance conversationally implicates that I can cook, and whoever understands it is aware that what I imply (that I can cook) 'follows from' what I say or my saying of it; i.e., whoever fully understands the utterance is aware of what is said, of what is implied, and of the inferential connection between what is implied and (the saying of) what is said.

The second feature of conversational implicatures I want to draw attention to is the global, post-propositional character of implicatures. Implicatures are generated via an inference *whose input is the fact that the speaker has said that p*. Hence no implicature can be computed unless something has been said, some proposition expressed. In particular, no implicature can be computed at a sub-locutionary level. We have to compute the truth-conditions first, so as to ascribe a definite content to the speaker's speech act, before we can infer anything from that speech act.

Grice's theory of conversational implicatures has enjoyed a tremendous success since he first put it forward in the sixties; but an important change has occurred, initiated by Grice himself. The notion of implicature has been extended to cases in which neither availability nor globality is present. Those are, indeed, the cases that matter most to semantics.

Consider the following examples :

- (1) Bill and Jane got married and had many children
- (2) Bill and Jane have three children

Sentence (1) implies that Bill and Jane got married before having the children. This is standardly accounted for by saying that the speaker is expected to respect the maxim of manner, which enjoins one to be orderly and, in reporting events, to report them in

the order in which they occurred.² Given that assumption, the speaker's saying (1) implies that the marriage took place before the birth of the children. Sentence (2) similarly implies that Bill and Jane have at most three children, for if they had more than three children the speaker ought to have said so (in virtue of the maxim of quantity which he is presumed to respect). So both the upper-bounded reading of the numeral in (2) and the temporal reading of the conjunction in (1) are said to result from enriching the core meaning of the sentence with a conversational implicature. This is a typical use of the notion of conversational implicature in contemporary discussions. What I find striking, however, is the lack of the two features I mentioned earlier. First, the discourse participants are not aware that the alleged implicature is not part of literal content ; the availability condition is not satisfied, in contrast to what happens in the 'I am French' type of case. Second, the alleged implicatures seem to occur locally. They fall within the scope of operators, as in the following examples :

- (3) Bill and Jane have three or four children.
- (4) Every father feels happy if his daughter gets married and gives birth to a child ; much less if she gives birth to a child and gets married.

In (3) the numerals are given the upper-bounded reading (*exactly three or exactly four*), but this can hardly be derived inferentially from the speaker's saying (3) in the minimal sense (*at least three or at least four*). The strengthening seems to occur locally, within the scope of disjunction, rather than globally. And in (4), the temporal suggestion is integral to the antecedent of the conditional. This contrasts with the normal behaviour of implicatures, which do not fall within the scope of operators because they arise at the speech act level, not at the level of sub-locutionary constituents.

Faced with this abnormal behaviour, one has to make room for a special class of implicatures (or pseudo-implicatures) which are intuitively undistinguishable from

² There is an alternative analysis, which does not rest on manner-based implicatures but on the fact that the tense features must be contextually assigned temporal values. (And there are other alternative analyses in the same vein.)

semantic content and can arise locally.³ Some theorists have appealed to the notion of a 'generalized' conversational implicature to handle such cases. Others have appealed to the distinction between genuine implicatures, which are distinct from and additional to what is said, and what Kent Bach has called 'conversational implicatures', i.e. things that are 'implicit *in* what is said' rather than implied by the act of saying it.⁴ My aim in this paper is to compare these (and other) approaches to the problem raised by what I will henceforth call 'embedded implicatures': seeming implicatures that arise locally, at a sub-locutionary level, without resulting from an inference in the narrow sense.

II. Generalized Conversational Implicatures: Two Conceptions

According to Grice, some conversational implicatures are 'generalized', i.e. they do not arise 'in virtue of special features of the context', but are normally carried by saying a certain thing or type of thing. The implicature arises 'in the absence of special circumstances', he says (Grice 1989: 37). The fact that, in a narrative, a conjunction such as 'They got married and had many children' is interpreted as mirroring the temporal order of the reported events is seen by Grice as resulting from a generalized conversational implicature: such an implicature is normally carried by an event-reporting conjunctive utterance such as (1).

The idea that some implicatures are generalized goes some way toward explaining why, in such cases, the availability condition is not satisfied. As Levinson puts it, generalized conversational implicatures are 'hard to distinguish from the *semantic* content of linguistic expressions, because [they are] routinely associated with linguistic expressions in all ordinary contexts' (Levinson 1983: 127). This explanation is not without its problems, but let it pass and consider how, using the notion of a generalized conversational implicature, we can account for the second of

³ I assume that the two properties go together, but this is only a conjecture. In this paper, however, I am concerned mainly with the second property (locality).

⁴ Bach 1994. In relevance-theoretic terminology (Sperber and Wilson 1986, Carston 2002), such things are said to be implicit constituents of the 'explicature', rather than implicatures.

the two facts noted above: the fact that the alleged implicature can arise at a sub-sentential level, as in (4). According to Mitchell Green, the connection is straightforward — a generalized implicature is *eo ipso* embeddable:

If assertion of a sentence *S* conveys the implicatum that *p* with nearly universal regularity, then when *S* is embedded the content that is usually understood to be embedded for semantic purposes is the proposition (*S* & *p*). (Green 1998: 77)

But why is that so exactly? Why is a generalized implicature — or at least, one that is 'nearly universal' — supposed to be embeddable?

Generalized conversational implicatures are still conversational implicatures, for Grice. To calculate an implicature, whether generalized or particularized, 'is to calculate what has to be supposed in order to preserve the supposition that the Cooperative Principle is being observed' (Grice 1989: 39-40): The implicatures are inferred from the speaker's saying that *p* and the presumption that he is observing the Cooperative Principle. The only difference between generalized and particularized implicatures lies in the amount of contextual information needed to derive the implicature from the speaker's speech act. When an implicature is generalized, the inference goes through 'independently of information about particular contexts of utterance'. We know that, *in general*, someone who says that *p* respects the Cooperative Principle only if *q* — and therefore we can infer that *q* from his saying that *p*, without having to rely on specific information about the context of utterance. Thus understood the notion of a generalized conversational implicature (henceforth to be called a 'Gricean' generalized implicature or GGI) is a graded notion: an implicature is more or less generalized (or particularized) depending on the amount of information regarding the context of utterance that is necessary to derive the implicature. But the mechanism of the derivation is the same in all cases: the implicatures are inferred from the performance of the locutionary act (i.e. from the speaker's saying that *p*), given the presumption that he is observing the Cooperative Principle. (See Carston 2002 : 111, and the literature cited therein.)

In this Gricean framework, however, we cannot account for sub-sentential cases. If conversational implicatures are the pragmatic implications of a speech act, they cannot arise at a sub-locutionary level. This point was made most explicitly by

Ducrot. In the late sixties Ducrot had independently come up with a theory of implicatures very similar to Grice's (Ducrot 1969), and in the early seventies he got involved in an in-depth study of semantic scales (e.g. Ducrot 1973, chapter 13). However, contrary to Horn and Fauconnier, who were exploring the same territory, he and his co-author Anscombe resisted the straightforward application of Gricean ideas to scalar phenomena.⁵ The alleged scalar implicatures, they argued, cannot be genuine implicatures because they fall within the scope of logical operators. In 'John has either five or six children', the numerals take the upper-bounded reading ('John has either *exactly five* or *exactly six* children') rather than the minimal reading ('John has *at least five* or *at least six* children'), but this cannot be due to an implicature since the implicature in question would fall within the scope of the logical connective 'either... or'. This, according to Anscombe and Ducrot, is impossible, in virtue of the following argument:

- (a) Conversational implicatures are pragmatic consequences of an act of saying something.
- (b) An act of saying something can be performed only by means of a complete utterance, not by means of an unasserted clause such as a disjunct or the antecedent of a conditional.
- (c) Hence, no implicature can be generated at the sub-clause level, i.e. at the level of an unasserted clause such as a disjunct or the antecedent of a conditional.
- (d) To say that an implicature falls within the scope of a logical operator is to say that it is generated at the sub-clause level, viz. at the level of the clause on which the logical operator operates.
- (e) Hence, no implicature can fall within the scope of a logical operator.

It follows that in examples like (3) and (4), the alleged implicatures responsible for the temporal reading of the conjunction or for the upper-bounded reading of the numerals are not genuine implicatures; for they are not inferred from the speaker's speech act but are constitutive of the proposition that is the content of that act.

⁵ See their collection of papers (Anscombe and Ducrot 1983), and especially their reply to Fauconnier 1976 (Anscombe and Ducrot 1978), reprinted therein.

At this point, however, we may be tempted to change the framework and alter the characterization of generalized conversational implicatures. From Grice's idea that generalized implicatures arise in the absence of special circumstances, there is but a short step to the conclusion that they are *generated 'by default', i.e. blindly, as soon as the relevant form of words is encountered*. According to Horn, Gazdar, and especially Levinson, who took that step, generalized implicatures are *default implicatures* (DI) : they are generated automatically (without inference). They belong to the 'micropragmatic' rather than to the 'macropragmatic' level, in Robin Campbell's typology:

A macropragmatic process is one constituted by a sequence of explicit inferences governed by principles of rational cooperation. A micropragmatic process develops as a cryptic [= unconscious] and heuristic procedure which partially replaces some macropragmatic process and which defaults to it in the event of breakdown. (Campbell 1981: 101)

Generalized conversational implicatures, thus understood, are no longer inferred from the speaker's saying that *p* together with the presumption that the Cooperative Principle is being observed. They arise through a different mechanism: they are generated by default when the relevant linguistic trigger is encountered, *unless* something in the linguistic or extralinguistic context blocks the generation and 'defeats' the implicature.

On this view generalized conversational implicatures are not merely generalized, they are also conventionalized: they are *associated with certain linguistic items* serving as triggers for the automatic process of implicature generation. That generalized implicatures tend to become conventionalized in this way seems natural. The conventions associating linguistic forms with DI arguably belong to the category of 'conventions of use', as opposed to straightforward 'meaning conventions' (Searle 1975, Morgan 1978). They are similar to the conventions in virtue of which an instance of the construction 'Can you VP?' is readily interpreted as a request, even though literally it is a question. The derivation of the indirect speech act of request from the direct speech act of question is based upon rationality considerations of the Gricean sort, but the inference is short-circuited as a result of generalization and conventionalization (Bach and Harnish 1979).

Levinson has explicitly resisted the equation of DI to 'standardized' or 'short-circuited' implicatures. The latter rely on compression by precedent and arise from routinization, he points out ; while default implicatures 'are generative, driven by general heuristics and are not dependent upon routinization' (Levinson 2000 : 24). This contrast may perhaps be interpreted as follows. Consider scalar implicatures (the paradigm case of DI). They are triggered by a specific form of words only because (a) that form of words (e.g. 'some') is conventionally recognized as belonging to a scale (e.g. the scale <some, most, all>), and (b) there is a meta-rule that determines, for every pair <S, *i*> consisting of a scale S and an item *i* belonging to that scale, the DI associated with *i*. In the case of short-circuited implicatures, arguably, the implicature is directly associated with a particular form of words through routinization, without there being any 'meta-rule' from which the conventional association flows. Let us assume that this makes sense and is what Levinson has in mind. Despite this alleged difference between short-circuited implicatures and DI, Levinson acknowledges that both belong to an intermediate layer between sentence meaning and speaker's meaning :

According to the standard line (more often presupposed than justified) there are just two levels to a theory of communication : a level of sentence-meaning (to be explicated by the theory of grammar in the large sense) and a level of speaker-meaning (to be explicated by a theory of pragmatics, perhaps centrally employing Grice's notion of meaning_{nn})... Speaker-meaning, or utterance-token-meaning, will be a matter of the actual nonce or once-off inferences made in actual contexts by actual recipients with all of their rich particularities. This view, although parsimonious, is surely inadequate, indeed potentially pernicious, because it underestimates the regularity, recurrence, and systematicity of many kinds of pragmatic inferences. What it omits is a third layer (...) of systematic pragmatic inference based *not* on direct computations about speaker-intentions but rather on general expectations about how language is normally used. (Levinson 2000 : 22)

The third layer involves conventions of use, in virtue of which certain forms of words, for one reason or another, come to be (defeasibly) associated with certain meanings over and above the meanings that are encoded at the first level.

To sum up, we must draw a distinction between two sorts of generalized implicatures. Gricean generalized implicatures (GGI) are still conversational implicatures and, as Ducrot and others (e.g. Cohen 1971) pointed out, they cannot arise at the sub-sentential level. Default implicatures (DI) are conventionally associated with certain linguistic forms, serving as triggers. Since they arise automatically rather than through Gricean reasoning, they can be embedded.⁶

The difference between DI and GGI is not merely the fact that DI are conventionally associated with certain forms of words. GGI themselves may get conventionalized, without ceasing to be conversational implicatures in the strict, Gricean sense. At a certain point, however, they will inevitably become DI. The transition from GGI to DI takes place when, as a result of conventionalization, a generalized implicature loses the property of 'nondetachability' which Grice uses to characterize conversational implicatures.

When saying a certain thing carries a conversational implicature, Grice claims, it is not possible to find another way of saying the same thing — another form of words, with the same content — which does not also carry the implicature. This is nondetachability. Now the existence of linguistic triggers for an implicature does not, *by itself*, preclude its being nondetachable. If asserting a disjunctive proposition carries the generalized implicature that the speaker is not in a position to separately assert any of the disjuncts, that implicature will be carried whichever form of words is selected to express the disjunctive proposition in question, even if, as the result of the generalization of the implicature, the word 'or', which is our primary means for expressing disjunction, has come to be associated with the implicature and acts as a 'trigger' for it. Still, *once a certain degree of conventionalization has been reached, a*

⁶ Because DI-theorists such as Levinson use the Gricean label 'Generalized Conversational Implicature', the notion of DI is commonly (and mistakenly) ascribed to Grice. Thus Bart Geurts writes: 'Grice's idea seems to have been that if a conversational implicature ϕ occurs often enough in the presence of an expression α , then the implicature will somehow become conventionally associated with α itself (Geurts 1998: 95-96). Note that the *Gricean* notion of generalized conversational implicature is immune to the criticisms Geurts addresses to the notion of DI and its use in the theory of scalars.

new possibility will arise. The 'implicature' will tend to be routinely generated even in configurations in which it could not result from a global inference à la Grice. That is how sub-sentential implicatures can be accounted for, in the revised framework. What starts life as a generalized implicature becomes conventionalized, and at a certain point is triggered even in contexts in which it could not be generated as an implicature via the Gricean post-propositional mechanism. At this point we no longer have a GGI, but a DI, characterized by the loss of the nondetachability feature. For consider the 'implicature' as it arises in a linguistic context (say, at the sub-sentential level) where it could not be generated via the Gricean mechanism. In such a context it arises only because there is a convention associating it to the form of words that happens to be used. Were it not for the convention of use in virtue of which the implicature is triggered by a certain form of words, it would not be generated, in such a context. It follows that, in such a context, the 'implicature' can be detached by changing the form of words that is used.⁷ It is no longer 'nondetachable'.

One might object that default implicatures, thus understood, are nothing other than Grice's conventional implicatures. But that is not so. What distinguishes conversational implicatures from conventional implicatures, according to Grice, are the two properties of nondetachability *and* cancellability. A conversational implicature can always be denied, cancelled or defeated, while this is not the case for conventional implicatures and other aspects of semantic content. Since default conversational implicatures are still cancellable (defeasible), they remain distinct from conventional implicatures, even though they do not possess a high degree of nondetachability.

III. From Pragmatics to Semantics

⁷ As Gazdar notes, 'to read off im-plicatures [i.e. default implicatures] from the semantic interpretation of the sentence (i.e., the proposition it expresses) would be impossible, since many different sentences can express a given proposition and many of these will not contain the scalar item and thus not carry the im-plicature' (Gazdar 1979 : 56).

Using the notion of 'default implicature' the two observations we started from can be accounted for. The default implicatures are not consciously available because they result from a 'cryptic and heuristic procedure', not from a macropragmatic inference conducted at the personal level; and they can arise locally because they are not generated through a global inference using as premiss the fact that the speaker has said that p , but are automatically triggered by certain expressions during the online processing of the utterance.

In recent work, several semanticists — most prominently Fred Landman (2000) and Gennaro Chierchia (2001) — have endorsed the notion of a default implicature (as opposed to the classical, Gricean notion of an implicature). They have put forward detailed proposals regarding the paradigm case : scalar implicatures. Both Chierchia and Landman reject what Landman calls 'the Gricean Root': the idea that the scalar operation that derives the implicature operates on the output of the grammar, where the output of the grammar is the proposition expressed by the complete utterance. Instead, they hold that the default implicature (or at least, the 'core' of the implicature) 'is derived at the earliest level in the grammatical derivation of the sentence asserted where an appropriate scale is available' (Landman 2000: 229). As Chierchia puts it, 'implicatures are not computed after truth-conditions of (root) sentences have been figured out; they are computed phrase by phrase in tandem with truth-conditions (or whatever compositional semantics computes)' (Chierchia 2001: 1). They are 'introduced locally and projected upwards in a way that mirrors the standard semantic recursion' (id.).⁸

In Landman's framework, a numeral (or any other scalar term) is a 'scalar trigger', i.e. it is associated with a scale of alternatives that is exploited in generating default implicatures. The implicatures (or rather, their 'core') are built from that scale, *as soon as possible in the grammatical derivation of the sentence*, by negating the items stronger on the scale. Consider, for example, sentence (5a):

(5a) Bill believes that there were four boys at the party.

We first derive the implicature-core (5c) at the level of the embedded sentence (5b):

⁸ The idea that implicatures may be computed at the phrasal level makes its first explicit appearance in Cornulier 1984: 663-4 (see also p. 689).

(5b) There were four boys at the party

(5c) They weren't more than four boys at the party

From there on, while compositionally building up the meaning of the complex sentence (5a), we build up simultaneously its implicature from the implicature-core (5c), following the semantic composition of (5a). Thus, from the level where the core of the implicature is derived, we successively build up the following pairs:

that there were four boys at the party

that there weren't more than four boys at the party

believe that there were four boys at the party

believe that there weren't more than four boys at the party

Bill believes that there were four boys at the party

Bill believes that there weren't more than four boys at the party

The last sentence corresponds to a default implicature of (5a) that cannot be generated in the classical, Gricean framework. We cannot generate the scalar implicature 'Bill believes that there weren't more than four boys at the party' by negating a piece of information stronger than (5a) on some scale. The same thing holds for (6a), which, by default, implicates (6b):

(6a) Every boy kissed three girls

(6b) Every boy kissed not more than three girls

As Landman points out, the global method leads us nowhere in a case like this. We cannot generate (6b) by negating a piece of information stronger than (6a) on some scale. To account for the scalar implicature (6b) we must give up the 'Gricean Root' and assume that the core of the implicature is derived *before* the universal quantifier comes into play. At an early level in the compositional process we derive the pair

x_n kissed three girls

x_n kissed not more than three girls

At a subsequent stage of the derivation *every boy* is introduced and we get:

(6a) **For every boy** x_n : x_n kissed three girls

(6b) **For every boy** x_n : x_n kissed not more than three girls

Of course, the implicatures thus generated by the computational system of grammar remain implicatures; they can be defeated or cancelled, by all sorts of means.

In Chierchia's framework, scalar terms and the complex expressions that contain them are associated with two meanings: the plain meaning of the expression, which is computed in the usual way, and its strengthened (upper-bounded) meaning which incorporates the scalar implicature. By default, the strengthened meaning is preferred; but the implicatures may be cancelled by the linguistic or extralinguistic context, in which case one falls back on the plain meaning. As in Landman's framework, the scalar implicatures, generated by negating the items stronger on the associated scale (or rather, the weakest of those items), are automatically introduced by the computational system of grammar, and their introduction takes place as soon as possible after a scalar term enters the computation. As composition proceeds, however, the implicatures that have been locally introduced can be filtered out. The originality of Chierchia's position lies in his suggestion that scalar implicatures are not only generated by default, but are also *removed by default* in certain linguistic contexts. The contexts in question are those that Fauconnier originally characterized as 'entailment reversing': negative sentences, antecedents of conditionals, and more generally downward-entailing environments. In such environments the plain meaning (without the implicature) becomes informationally stronger than the strengthened meaning (with the implicature), so that maintaining the implicature would lead to a weakening of information content. The default generation-and-removal of scalar implicatures therefore mimicks, within grammar, the Gricean search for maximal informativeness. We may perhaps think of the Gricean post-propositional mechanism as being the evolutionary source of the grammatical mechanism which Chierchia

describes. It is as if a pragmatic mechanism had been incorporated into the design of grammar to make it more efficient.⁹

Even though it presumably evolved from a pragmatic mechanism involving the Gricean maxim of quantity, the default generation of scalar implicatures is not itself a pragmatic mechanism in the full-blooded sense: as both Landman and Chierchia make clear, it belongs to the computational system of grammar. In this respect Chierchia's and Landman's proposals are similar to that put forward by Jonathan Cohen in his early assault on Grice. Shortly after Grice delivered the William James Lectures, Cohen criticized his 'Conversational Hypothesis' on the grounds that it cannot account for embedded implicatures; and he offered his own 'Semantical Hypothesis' as a viable alternative (Cohen 1971). The Conversational Hypothesis says that the implicatures associated with the logical connectives result from a conversational inference, while Cohen's alternative 'Semantical Hypothesis' says they are part of the meaning of the connectives. Cohen was well-aware of Grice's criterion, Modified Occam's Razor, which says that senses should not be multiplied without necessity. But he thought it was possible to ascribe a single, unequivocal meaning to the logical connectives, by treating certain aspects of that single meaning as cancellable. Thus both Grice and Cohen ascribe a single meaning to the connectives, in conformity to Modified Occam's Razor. According to Grice, that meaning can be contextually enriched as a result of a pragmatic inference; according to Cohen, it can be impoverished by cancelling out a defeasible aspect of that meaning. I understand Landman and Chierchia as putting forward an account which, like Cohen's, is based on the acceptance of *defeasible semantic features*. Such an account is semantic, not pragmatic.

At this point it is worth reflecting on what, in the overall process of interpreting an utterance, distinguishes the mechanisms or processes that are 'pragmatic' from those that are 'semantic'.

⁹ Levinson speculates that such a mechanism was indeed needed to overcome what he calls the 'encoding bottleneck'. 'The actual process of phonetic articulation', he says, 'is a bottleneck in a system that can otherwise run about four times faster' (Levinson 2000 : 6). The solution to the bottleneck is this : 'find a way to piggyback meaning of top of the meaning' (ibid.).

A paradigmatically pragmatic process or mechanism such as the Gricean generation of (classical) conversational implicatures possesses the following features :

- it appeals to extralinguistic information : facts regarding the situation of utterance or the ongoing conversation, background knowledge, etc.
- it is 'top down' rather than 'bottom-up'; that is, it is not triggered by something linguistic — some aspect of the linguistic signal being processed — but takes place in order to make sense of the communicative act performed by the speaker. (Note that this contrast is not the same as the previous one. A process may invoke extralinguistic — contextual — information while being linguistically triggered in a typically bottom-up manner. Indexical resolution is a case in point : the process of contextually assigning a value to an indexical is triggered by the occurrence of that indexical in the sentence, yet extralinguistic information is clearly and crucially involved.)
- it is global rather than local ; that is, it is not part of the stepwise process of compositionally determining a semantic interpretation for the sentence, but takes place after the global interpretation of the sentence has been calculated.
- it is transparent ('available') to the users of the language because it is a matter of 'speaker's meaning', and speaker's meaning is essentially overt in the sense glossed by Grice and his followers : the conversational protagonists must be consciously aware of what the speaker means, while they need not be consciously aware of the grammatical meaning of the expressions used, nor of the processes through which the meaning of the whole is determined on the basis of the meanings of the parts.
- the output it delivers enriches the interpretation of the utterance in an optional manner ; that is, there are contexts in which the same form of words would carry the plain, unenriched interpretation.

The default generation of scalar implicatures described by Chierchia and Landman (the 'DGSI', for short) possesses only the last of these features. This is not sufficient to make it a pragmatic process. Extralinguistic information plays no role — it only comes into play to defeat the default implicature or to reinstate it (to 'freeze' it) in case of default removal. The DGSI is clearly bottom-up : it is triggered by the

occurrence of scalar terms in the sentence. It takes place locally and subpersonally, as part of the compositional process of determining the (default) truth-conditions of the sentence. Those features, and especially the fact that it is linguistically triggered and automatic (context-independent), put the DGSi squarely on the semantic side, despite the optional (defeasible) character of the output. In this regard the DGSi is a bit like the process of indexical resolution. That process too possesses only one out of the five features which characterize paradigmatically pragmatic processes. As we have noticed in passing, the process of indexical resolution is linguistically triggered (bottom-up). It takes place locally and subpersonally in the derivation of the sentence's truth-conditions. It is mandatory rather than optional (i.e. we *have to* assign a contextual value to the indexical, in virtue of the rules of the language). The only thing that is pragmatic here is the fact that contextual, extralinguistic information is appealed to in assigning a value to an indexical.

The following table summarizes the similarities and contrasts between the three processes we have been talking about : the Gricean post-propositional mechanism (GPM) construed as paradigmatically pragmatic, the default calculation of scalar implicatures described by Landman and Chierchia (DGSi), and indexical resolution (IR).

	<i>Extralinguistic information?</i>	<i>Personal-level Availability?</i>	<i>Global?</i>	<i>Top-down?</i>	<i>Optional?</i>
GPM	yes	yes	yes	yes	yes
DGSi	no	no	no	no	yes
IR	yes	no	no	no	no

IV. Pragmatic implications of sub-locutionary acts?

Even though the classical Gricean approach cannot handle embedded implicatures, while the semantic approach can, one may still attempt to account for them in a pragmatic (rather than in a semantic) framework, by giving up some aspect of the Gricean picture. For example, one may construe the relevant implicatures as pragmatic implications of something other than a self-standing speech act.

Recall the anti-Grice argument put forward by both Ducrot and Cohen:

- (a) Conversational implicatures are pragmatic consequences of an act of saying something.
- (b) An act of saying something can be performed only by means of a complete utterance, not by means of an unasserted clause such as a disjunct or the antecedent of a conditional.
- (c) Hence, no implicature can be generated at the sub-locutionary level, i.e. at the level of an unasserted clause such as a disjunct or the antecedent of a conditional.

It is possible to reject (b) by giving a weaker interpretation of the notion of 'saying'. On that interpretation, one 'says' (though one does not 'assert') something by uttering a disjunct or the antecedent of a conditional, and one's saying what one says may carry conversational implicatures. Alternatively, if one sticks to the strong interpretation of the notion of 'saying', according to which one does not 'say' anything by uttering a disjunct or the antecedent of a conditional — if, therefore, one accepts (b) — then one may reject (a) and claim that conversational implicatures need not be pragmatic implications of an act of *saying* (in the strong sense) but may be also pragmatic implications of an act of 'representing' or 'describing' — where representing or describing are things that *can* be done by means of an unasserted clause such as a disjunct or the antecedent of a conditional.

The line I have just described (rejecting [a] or [b], depending on the interpretation that is given of the notion of 'saying') is basically that taken by Ralph Walker in his reply to Cohen's criticism of Grice (Walker 1975). Walker argues that the unasserted antecedent of a conditional is nevertheless *uttered* when the conditional is uttered; and an implicature, he says, is a pragmatic implication of an utterance act — not necessarily of a full-blooded illocutionary or even locutionary act. He writes:

[The Conversational Hypothesis] holds that by a particular utterance on a particular occasion the speaker can convey more than his utterance strictly means through relying on a general recognition of Grice's Co-operative Principle. It is therefore concerned with utterances, whether they constitute self-standing speech acts or not; an utterance of a subordinate clause, as in the antecedent of a conditional, is still an utterance, and therefore may convey conversationally more than it literally means. It may convey, for example, a

further condition on which the consequent is to be taken to depend. (Walker 1975: 151)

Consider Cohen's original example:

- (8) If the old King has died of a heart attack and a Republic has been declared, then Tom will be quite content.

It involves an implicature of temporal order standardly accounted by appealing to the sub-maxim of orderliness, but one that, in this particular example, arises locally at the level of the antecedent of the conditional. (Tom will be content, on Cohen's scenario, only if the King died *before* a Republic was declared.) If Walker is right, this should not be a problem. The maxim of orderliness arguably demands that, in representing or describing sequences of events (whether assertively or not), one refrain from representing them in a different order than the order in which the speaker wants the addressee to think of them as having taken place (whether the addressee's 'thinking' itself is assertive thinking or mere entertaining). If the speaker wants the addressee to think of A as having taken place before B, he should, in virtue of the maxim, frame his description in such a way that the representation of A precedes the representation of B, that is, in such a way that, *in the discourse*, A is introduced before B. In this way the speaker spares the hearer unnecessary efforts. Now this constraint is a constraint on how temporally ordered events are represented or described; and such representation/description of sequences of events may surely take place in the antecedent of conditionals, or in unasserted clauses more generally. Thus in (8) the antecedent contains a description of two events: the King's death, and the declaration of a Republic. The speaker does not assert that those events took place: he speaks conditionally. But the events are nonetheless described, and the sub-maxim of orderliness therefore applies. In virtue of the presumption that the speaker respects the maxims, his describing the two events in a certain order suggests that, when making the supposition corresponding to the antecedent of the conditional, the hearer is to think of those events as having taken place in that order. (Or at least, this suggestion will be conveyed if the temporal order of the events is communicably relevant.) Uttering the antecedent therefore carries a conversational implicature which enriches the content of the supposition the hearer is

invited to make; the implicature provides, as Walker puts it, 'a further condition on which the consequent is to be taken to depend'. Or consider a belief report such as 'Paul believes that the King has died and that a Republic has been declared'. The speaker's describing the two events in that order suggests that, according to Paul, they took place in that order. The speaker does not assert that they took place in that order (not even that they took place): he reports Paul's beliefs. Still, his describing the events in a certain order in reporting Paul's beliefs carries an implicature relative to the temporal ordering of the events in Paul's beliefs.

In this way, with a little effort (in order to make the proposal more precise), we can account for some of the problematic cases. But it is not certain that the strategy I have outlined (following Walker) can be generalized and account for all the cases. In particular, it is unclear how it would apply to scalar implicatures. The scalar reasoning appeals to the idea that the speaker respects the maxim of quantity, i.e. gives as much (relevant) information as possible; now it is far from obvious that the notion of 'giving information' can be divorced from that of asserting (or from similar notions), as the strategy requires.

In some cases, admittedly, the strategy *can* be invoked in dealing with alleged scalar implicatures. For example, when the scalar term receives focal stress, the implicature may be construed as arising as a pragmatic implication not of the act of asserting or giving information, but of the sub-locutionary act of *stressing a particular word* (something that may happen in an embedded clause). By putting focal stress on a word, one implicates that the alternatives to that word in a contextually salient contrast set do not apply, i.e., that the result of substituting them for the word bearing focal stress would not be acceptable.¹⁰ If the word is a scalar term, the salient alternatives will typically be the items on the associated scale, or perhaps the items on the upper part of the associated scale. The resulting implicature will therefore look very much like a scalar implicature, even if the mechanism through which it is generated is quite different.¹¹ Be that as it may, the strategy I have described,

¹⁰ This is a simplification. The expression in focus need not be the word actually bearing stress but may be a larger constituent containing it.

¹¹ I think the implicatures generated by putting focal stress on a word — whether scalar or not — are best treated as *conventional* implicatures.

following Walker, unproblematically applies to scalar implicatures only in that sort of case. It is unclear that it can account for the cases in which a scalar implicature is in the scope of an operator without the scalar word being stressed.

V. 'Local' speech acts?

Alternatively, we can maintain that conversational implicatures are pragmatic implications of a full-blooded speech act, while rejecting the claim that no such speech act is performed by uttering a disjunct or the antecedent of a conditional. On this view the sentential parts of a complex sentence are used to perform speech acts of their own. Just as the complex sentence is built up from its parts, the speech act it serves to perform is also built up from the local speech acts which the sentential parts of the complex sentence are used to perform.

The best example of that sort of move is provided by the speech-act theoretic analyses of conditionals put forward in the early seventies (see e.g. Ducrot 1972, Mackie 1973). According to these analyses, a conditional sentence 'If P, then Q' may be construed as serving to perform a complex speech act consisting of two local speech acts: a first speech act whereby the speaker makes a supposition, and a second speech act whereby, in the scope of that supposition, the speaker asserts something. The first speech act is performed by uttering 'If P', and the second speech act by uttering the consequent in the context created by the first speech act. In this framework, nothing prevents the first speech act, performed by uttering the antecedent, from carrying conversational implicatures of the standard sort.

This strategy is that which has been used by Stalnaker to solve the projection problem for presuppositions within a pragmatic framework (Stalnaker 1974; see also Karttunen 1974). Consider a conjunctive statement 'P and Q', where Q presupposes R. Although the presuppositions of the parts are normally inherited by the whole, there are cases in which the complex sentence 'P and Q' will not presuppose R — for example if P itself entails R (Karttunen 1973). How can we account for that fact? Stalnaker offers the following explanation. The speaker who says 'P and Q' *first asserts P and then asserts Q*. For Q to presuppose R is for it to be assertable *only* in a context in which R is assumed to hold. Now the speaker's asserting P changes the context by adding P to the common ground, in such a way that the consequences of P, including R, will themselves be part of the common ground when the consequent,

Q, is uttered. It follows that 'P and Q' will be assertable not only in contexts in which R is antecedently assumed to hold, but in any context (since the presupposition normally carried by Q is 'internally' satisfied by the first conjunct in the complex sentence 'P and Q'): hence 'P and Q', contrary to 'Q', does not presuppose that R (i.e. it is not 'assertable only in contexts in which R is assumed to hold').

Stalnaker gives the same explanation for conditional sentences in which a presupposition of the consequent is internally satisfied by the antecedent, as in 'If France has a king, the king of France is bald'. Here too he describes the discourse as involving two speech acts: a first speech act of supposition, and an act of assertion performed in the temporary context created by the first speech act. The same sort of analysis can easily be extended to disjunctions such as 'Either France is a Republic, or the King of France is so shy that one never sees him in public'.

We could adapt this analysis to our examples involving embedded implicatures. Faced with a difficult case such as 'John has five or six children', where the *exactly*-implicatures fall within the scope of the disjunction, we may argue that, just as 'P and Q' consists of two successive assertions in the Stalnakerian framework, 'P or Q' also consists of two successive assertions. The difference between 'P and Q' and 'P or Q' is that the first sequence of assertions is conjunctive, while the second sequence is disjunctive. To say that a sequence of two assertions is disjunctive is to say that the second assertion qualifies the first, and is presented as holding just in case the first assertion turns out to be false. 'P or Q' therefore means something like: 'P; but if not-P, then Q'.

It is easy to check that, on such an analysis, we can handle embedded implicatures. The speaker first asserts that John has five children, thereby conveying the implicature that he has no more than five. Then, by saying 'Or he has six', he asserts that, if John does not have exactly five children, he has six (thereby conveying the implicature 'no more than six').

The problem with this analysis is that it blurs an intuitive distinction between a disjunctive sequence of assertions, henceforth to be called a 'disjunctive assertion', and the assertion of a disjunction; a distinction analogous to that between a conditional assertion (e.g. 'If you are hungry there are cookies in the sideboard') and the assertion of a conditional. As an example of disjunctive assertion, consider:

John has five children. Or he has six.

John has five children; or he has six.

John has five children, or six.

Here, clearly, a first assertion to the effect that John has five children is followed by a second assertion, introduced by 'or'. 'Or' indicates that the second assertion is an alternative to the first assertion. This, Cornulier suggested (1982: 88-90), can be cashed out by representing the content of the second assertion as a conditional whose antecedent is the negation of the first assertion (i.e. the negation of what it asserts).¹² So far so good. But when I say 'John has five or six children', it does not seem that I first assert that John has five children and then assert that in the opposite case he has six. This example is most naturally understood as the assertion of a (single) disjunctive proposition — a reading which can be made explicit by using 'either...or...': 'Either John has five children or he has six'. This *cannot* be interpreted as a disjunction of assertions; for the speaker at no point asserts that John has five children: from the very start, the proposition that John has five children is presented as one of the disjuncts, only the disjunction being asserted. Yet this does not prevent the embedded *exactly*-implicature from arising.

In response to this objection, one might grant the intuitive distinction between a disjunctive assertion and the assertion of a disjunction (or between a conditional assertion and the assertion of a conditional), while holding that disjunctions, like conditionals, can themselves be analysed in speech-act theoretic terms. It is true that, when we assert a disjunction, we do not separately assert either the disjuncts. But the consequent of a conditional is not *really* asserted either, and that fact does not rule out a speech-act-theoretic analysis according to which the consequent is *locally* asserted, i.e. asserted in the local context set up by the supposition of the antecedent. The same sort of analysis in terms of local speech acts and local contexts might be attempted for disjunctions, and for complex utterances generally.

But this will not do. If we treat any of the disjuncts in 'Either John has five children or he has six' as locally asserted, as suggested, then it is clear that the notion of local assertion we use is not the full-blooded speech-act-theoretic notion of assertion, but a semantic surrogate. This is something that has been independently

¹² Cornulier's suggestion concerns disjunctive (or, as he put it, 'alternative') *questions*, but his proposal easily generalizes, as he himself points out (Cornulier 1982: 99-101).

noted in connection with the Stalnakerian analysis of conditionals. Like Ducrot and Mackie, Stalnaker and his followers say that when we assert a conditional, we assert the consequent in the local context created by uttering the antecedent. As Landman emphasised, however, the 'local context' in which that assertion takes place is not a *real* context:

The context in which we evaluate the assertion of the consequence is not the actual speech context, but a context which derives from the actual speech context by adding the antecedent. This context is called the *local context*. But of course the consequent *isn't* asserted in the local context, and the local context *isn't* an actual speech context. (...) The presuppositions that derive from the actual assertion of the sentence in the actual speech context are characterized in terms of what the parts of that sentence *would* presuppose if they *were* asserted in a local context.¹³ (...) The local context is derived from the actual speech context, following the semantic composition of the sentence. This means that the notion of local context is a *grammatical* notion. (Landman 2000: 237)

Landman then goes on to stress the difference between the pragmatic notion of context standardly used in implicature theory, and the grammatical notion of local context used in presupposition theory:

While presupposition theory has been regarded since the early seventies as basically a theory of local context, the notion has been all but ignored in implicature theory... The reason is, I think, that unlike for presuppositions, there is a consensus that *implicatures are derived directly from the actual speech situation*. If local context is relevant for implicature as well, this means that we cannot, following the Gricean Root, let the theory of co-operative information exchange derive them from the actual assertion of the sentence in the actual context. It means that we have to let the grammar derive them from implicatures that would be derived if its parts were asserted in a local context,

¹³ The emphasis in this sentence is mine (FR).

which itself is derived from the actual context following the semantic composition. (Landman 2000: 237)

I conclude that, suitably weakened so as to avoid the confusion between a disjunctive (or conditional) assertion and the assertion of a disjunction (or conditional), the alleged speech-act-theoretic analysis inspired by Stalnaker collapses into a semantic analysis of the sort advocated by Chierchia and Landman. It is not a genuine alternative to that sort of analysis.

VI. Embedded implicatures and 'free enrichment'

A last strategy is available to the theorist who wants to resist the semanticization of embedded implicatures. It relies on the distinction between primary and secondary pragmatic processes (Recanati 1989, 1993, 2001, 2004). Secondary pragmatic processes are post-propositional inferences à la Grice: in interpreting an utterance, what is implied, in the intuitive sense, is inferentially derived from the speaker's saying what s/he says (in a way that satisfies the availability condition). In contrast, primary pragmatic processes are pragmatic processes at work in the very determination of what is said. For example, we need to assign indexicals and other context-sensitive expressions a contextual value in order to fix truth-conditional content. This contextual process of value assignment, which I call 'saturation', is irreducibly pragmatic: considerations relative to what the speaker means play a crucial role in that process. Still, it contributes to the determination of the utterance's semantic (truth-conditional) content, hence it counts as a 'primary' pragmatic process.

In contrast to secondary pragmatic processes, primary pragmatic processes operate locally rather than globally, and they do not satisfy the availability condition. In these respects they are like the default generation of implicatures. Qua pragmatic processes, however, they tap extralinguistic information and appeal to nonlinguistic abilities like the ability to decipher intentions and to make sense of actions.

The only primary pragmatic process that is standardly acknowledged is the process of saturation in virtue of which indexicals and other context-sensitive expressions are assigned a contextual value. As we have seen, saturation is a bottom-up process (i.e. it is triggered by some element in the sentence) and it is mandatory (a value *must* be contextually provided). On the standard picture, any

other pragmatic process involved in interpreting an utterance is considered as secondary, i.e. post-propositional, in the manner of the GPM. This picture rests on the idea that there is pragmatics on one side and truth-conditions on the other side. The reason why the contribution of pragmatics to truth-conditions is allowed for in the case of indexicals is that the pragmatic process at stake is triggered by something linguistic — hence it's not purely contextual — and it is mandatory rather than optional, so that it cannot be dispensed with anyway. But that is the only exception that is made to the principle that pragmatics has no bearing on truth-conditions. As Stanley puts it,

All effects of extra-linguistic context on the truth-conditions of an assertion are traceable to elements [e.g. indexicals or free variables] in the actual syntactic structure of the sentence uttered. (Stanley 2000: 391)

But the standard picture which this quotation illustrates has been questioned. Apart from the desire to keep pragmatics away from the business of determining truth-conditions, there is no good reason to deny the existence or at least the possibility of primary pragmatic processes that, unlike saturation, are not linguistically but contextually triggered and are optional rather than mandatory. That there are such processes is the gist of the alternative picture known as 'Truth-Conditional Pragmatics' (TCP).

According to TCP, saturation is *not* the only pragmatic process that is primary and can affect truth-conditional content. Consider, for example, the process which Geoff Nunberg has dubbed 'predicate transfer' (Nunberg 1995). It takes us from a certain property, conventionally expressed by some predicative expression, to a distinct property bearing a systematic relation to it. For example, in 'I am parked out back', 'parked out back' undergoes predicate transfer. The property that is literally encoded is a property of cars (the property of being parked out back), but the property which the expression actually contributes to the truth-conditions in this utterance is not a property of cars but another, systematically related property, namely the property a car-owner has when his or her car has the former property.

In contrast to saturation, the process of predicate transfer is neither mandatory nor bottom up. It may look as if, in an utterance such as 'I am parked out back', that process *must* take place, because there is a linguistic mismatch between the

predicate (which denotes a property of cars) and what it is applied to (a person). But type-mismatch is not necessary for predicate transfer. Just as, through transfer, 'The ham sandwich left without paying' is understood as saying something about the customer who ordered the sandwich, 'The ham sandwich stinks' can be so understood, in a suitable context, even though the property of stinking potentially applies to sandwiches as well as to customers.¹⁴ The process of transfer is not a linguistically controlled but a pragmatically controlled pragmatic process: it is not triggered by something linguistic — some aspect of the linguistic signal being processed — but takes place in order to make sense of the communicative act performed by the speaker. Moreover, it is optional: there are contexts in which the same form of words would carry the plain interpretation, without transfer. In some contexts, 'The ham sandwich stinks' talks about the sandwich; in other contexts, through transfer, it talks about the customer. Whether or not predicate transfer takes place is a wholly pragmatic matter. It is not something that is dictated by linguistic conventions.

Even though the pragmatic process of predicate transfer is optional (rather than mandatory) and top-down (rather than bottom up), still it takes place locally and interferes with the process of semantic composition. In another classic example, 'There is a lion in the courtyard', 'lion' can be understood, through transfer, in the representational sense: the thing that is said to be in the courtyard is not a (real) lion but a *representation* (more specifically, a statue) of lion. Now consider 'There is a stone lion in the courtyard'. What is said to be made of stone here? Clearly, it is the statue, rather than the lion which the statue represents. This simple fact shows that the process of representational transfer which affects the word 'lion' must take place *before* the composition rule associated with the noun-noun construction applies to the semantic values of the nouns 'stone' and 'lion'.¹⁵ If predicate transfer applied

¹⁴ I am indebted to Dan Sperber for this example. The original ham-sandwich example is, of course, Nunberg's.

¹⁵ Note that this composition rule itself is context-sensitive (Partee 1984: 294-5). The denotation of the compound results from intersecting the (literal, or pragmatically derived) denotation of the head noun with the set of objects that bear a certain relation R to the (literal, or pragmatically derived) denotation of the modifying noun. That

globally, after the grammatically triggered composition rules have applied, the interpretation we would get for the noun-phrase 'a stone lion' would be something like: *a representation of (a lion that is made of stone)*. But the correct interpretation is: *(a representation of a lion) that is made of stone*. We must therefore give up the Gricean idea that pragmatic processes operate globally on the output of the grammar.¹⁶

According to TCP, just as indexical expressions are standardly assigned two levels of semantic value ('character' and 'content'), we must distinguish two levels of semantic values for ordinary, nonindexical expressions. Semantics assigns lexical values to simple expressions; pragmatics then optionally comes into play to determine the *compositional values* which those expressions assume in the linguistic and extralinguistic context in which they occur. Compositional values, not lexical values, are what undergo semantic composition. In other words, the composition rules determine the value of the whole on the basis of the *pragmatically determined* compositional values of the parts.

Predicate transfer is only one among a family of pragmatic processes that have the properties I have listed: they are pragmatically controlled (top-down) rather than linguistically controlled (bottom-up), they are optional rather than mandatory, and they take place locally, thereby interacting with the compositional determination of truth-conditional content. This family of primary pragmatic processes I call 'modulation', as opposed to saturation (Recanati 2004). Modulation takes as input the meaning of some expression (whether simple or complex) and returns as output a pragmatically derived meaning serving as compositional value.

Among the processes of modulation that affect the truth-conditions of utterances, the most typical and pervasive is free enrichment,¹⁷ in virtue of which an expression is contextually given a more specific interpretation than it literally

relation can only be contextually determined. In 'stone lion', 'R' is typically assigned the relation *being made of*, but in less accessible contexts a different relation will be assigned to the variable.

¹⁶ See Sag 1981 and Jackendoff 1997: 55 and 65-66 for similar points.

¹⁷ In the pragmatics literature this process is also called 'narrowing', 'strengthening', or 'expansion'. (There are a couple of other labels as well.)

encodes. Through free enrichment an expression takes a pragmatically derived denotation that is a subset of the initial denotation (Carston 1997). Thus, to take another classic example (discussed by Nunberg and Zaenen 1992) we understand the mass term 'rabbit', which literally means something like *rabbit stuff*, as meaning *rabbit fur* in 'She wears rabbit' and *rabbit meat* in 'She eats rabbit'. Or, to take an example discussed by Searle, we understand the word 'cut' differently in 'cut the grass' and in 'cut the cake':

Though the occurrence of the word "cut" is literal in [both] utterances..., and though the word is not ambiguous, it determines different sets of truth conditions for the different sentences. The sort of thing that constitutes cutting the grass is quite different from, e.g., the sort of thing that constitutes cutting a cake. One way to see this is to imagine what constitutes obeying the order to cut something. If someone tells me to cut the grass and I rush out and stab it with a knife, or if I am ordered to cut the cake and I run over it with a lawnmower, in each case I will have failed to obey the order. That is not what the speaker meant by his literal and serious utterance of the sentence. (Searle 1980: 222-223)

Both 'cut' and 'rabbit' are given contextually specific interpretations through free enrichment. Of course, the linguistic context plays an obvious role here, but free enrichment remains a contextually-driven (top-down) and optional process. Nothing prevents 'cut' in 'cut the grass' from being contextually interpreted in the sense of *slice into strips*, or 'rabbit' in 'She wears rabbit' from being interpreted in the sense of *rabbit meat*. And nothing prevents the meaning of either expression from remaining contextually plain and unenriched ('After the accident, there was rabbit all over the highway'). These properties, which free enrichment shares with the other pragmatic processes in the modulation family, are not exhibited by processes of the saturation family (indexical resolution, etc.). We must therefore complete our table and make room for modulation alongside the three types of process we have already described (the Gricean post-propositional mechanism, the default generation of scalar implicatures, and saturation):

	<i>Extra-linguistic information?</i>	<i>Personal-level Availability?</i>	<i>Global?</i>	<i>Top-down?</i>	<i>Optional?</i>
Grice's post-propositional mechanism	yes	yes	yes	yes	yes
Default generation of scalar implicatures	no	no	no	no	yes
Saturation (indexical resolution etc.)	yes	no	no	no	no
Modulation	yes	no	no	yes	yes

Once we acknowledge the primary pragmatic processes of modulation, a new approach to embedded implicatures becomes available. We can construe them as a particular case of free enrichment, whereby the meaning of e.g. scalar terms is contextually strengthened. This is, indeed, the view which many advocates of Truth-Conditional Pragmatics (e.g. Bach 1994, Bezuidenhout 2002) actually hold. They take the alleged implicatures to be 'implicatures' i.e. aspects of the proposition expressed which are provided by freely enriching or expanding the literal meaning of the sentence. The fact that the alleged implicatures fall within the scope of operators is taken to be the litmus test showing that they are not really conversational implicatures derived through the Gricean mechanism, but pragmatic constituents of what is said (Recanati 1989: 112-14, 1993: 269-74; Carston 2002: 191-7, forthcoming).

VII. Conclusion: Default Implicatures or Free Enrichment (or Both)?

We have seen that there are two viable approaches to embedded implicatures: a semantic approach in terms of default implicatures, and a pragmatic approach in terms of free enrichment. Which one is to be preferred? Well, I am not sure that we really have to choose. To conclude this paper, I will argue that the two views we have discussed do not necessarily stand in competition to each other.

It is true that both theories provide an account of embedded implicatures: both make room for a process of strengthening that is optional and takes place locally. It is true also that they offer conflicting characterizations of the process in question. According to one theory, that process is context-independent and belongs to the computational system of grammar. According to the other theory, it is a fully pragmatic, context-driven process. Still, I think there is no downright incompatibility

between these approaches, appearances notwithstanding. They are incompatible only if we assume that there is a *single* process at stake. But we cannot presuppose that that is so, for, as we shall see, this is one of the questions at issue in the debate.

DGSI-theorists argue that scalar strengthening is linguistically triggered and context-independent. But the default implicatures they posit *can* be defeated, in a suitable context (if they couldn't, they wouldn't be conversational implicatures). From the point of view of DGSI theory, the extralinguistic context plays a role at least in the *fixation* of the implicature — a contextual process which determines whether or not the default implicatures are defeated or (in the event of default removal) whether or not they are 'frozen' and maintained despite the downward entailing operator. It follows that there are two processes at work in the interpretation of scalar implicatures, and two components in the overall theory. One component belongs to semantics: it concerns the default generation of implicatures. The other component is pragmatic and concerns *what happens to the defaults when the sentence is uttered in a real context*. Since there are two components, and two processes jointly at work in the interpretation of scalar utterances, it may be that the two conflicting characterizations of 'the' local process of optional strengthening I have mentioned are in fact nonconflicting characterizations of two distinct processes. It may be that there is both a linguistically triggered, context-independent process of default generation of *potential* implicatures, and a pragmatic process of free enrichment taking as partial input the output of the previous process and leading to the *actual* strengthening of the meaning of the scalar term.

Of course, one need not accept the existence of default implicatures in the first place: one may deny their existence and account for embedded implicatures solely in terms of free enrichment. Granted; one need not accept the existence of free enrichment either. But nothing prevents a theorist from accepting *both* default implicatures and free enrichment — that is my point. At bottom, there are two distinct questions. If we provide a positive answer to either of the two questions, we have a solution to the problem of embedded implicatures. Anyone who gives a positive answer to one question may therefore safely give a negative answer to the other. But he or she does not *have to* give a negative answer to the other question. In other words, the two issues are orthogonal.

The first of the two orthogonal issues concerns defaults. Are there default pragmatic values (e.g. default scalar implicatures) whose calculation is part of the

computational system of language? Following suggestions by Gazdar and Levinson, Landman and Chierchia provide an affirmative answer to this question. Sperber and Wilson, their followers and a few other researchers (e.g. Geurts 1998) provide a negative answer. The second issue concerns pragmatics, and more specifically the context-driven, optional processes of modulation. Are such processes primary? Do they operate on the meanings of the parts before the meaning of the whole is calculated? Do they take place locally so as to affect semantic composition? Advocates of Truth-Conditional Pragmatics provide an affirmative answer to this question. Defenders of the standard, Gricean picture give a negative answer: pragmatic processes are essentially global and post-propositional, they say. Unless they are linguistically triggered and mandatory (as saturation is), they do not affect truth-conditional content.

To say that the two issues are orthogonal is to say that there are *four* possible positions, depending on one's answer to these questions. Let me review each of these positions in turn.

The first position, NN (for 'No-No'), accepts *neither* default pragmatic values nor primary pragmatic processes over and beyond saturation. It sticks to a classical conception of both semantic content (which does not include a defeasible layer) and pragmatics (which does not interfere with semantic composition).

How, if one takes such a position, can one account for embedded implicatures? One has to say that they do not really exist. Thus, according to King and Stanley, who defend the classical view, embedded implicatures are an illusion (King and Stanley, forthcoming, section V). Consider the following example:

(9) Eating *some* of the cake is better than eating all of it

Here a scalar implicature seems to enrich the left-hand-side of the 'better than' relation. (It is better to eat some-but-not-all of the cake than to eat all of it.) This is a typical case of embedded implicature. King and Stanley think such examples can be analysed in terms merely of saturation. A statement of the form 'Better P than Q' is true, they say, if and only if the most similar worlds in which the left-hand-side (P) holds are preferable (in some contextually determined sense) to the most similar worlds in which the right-hand-side (Q) holds. To be evaluated, such a statement

requires the contextual provision of a specific similarity relation between worlds.¹⁸ This, they think, is an instance of saturation. So they offer the following analysis, which dispenses with embedded implicatures. In (9) the word 'some' conveys its literal content (*at least some*), not its strengthened meaning (*some but not all*); but focussing 'some' leads the interpreter to choose a similarity relation such that the most similar P-worlds are worlds in which *not all* the cake has been eaten. In this way the truth-conditions of the utterance are affected as if 'some' had been given the upper-bounded reading (even though it has not). They conclude:

By focussing the relevant word, one affects the choice of the similarity relation between worlds that is relevant for the truth-conditions of the "better-than" construction in that context. So the truth-conditions of these constructions are affected by scalar facts, but independently of processes such as explicature or implicature "intrusion". Nor does the scalar information "enrich" the semantic content. Rather, the truth-conditions of "better-than" sentences are sensitive to the choice of a similarity relation between worlds, and focus affects the choice of that relation. (King and Stanley, forthcoming)

King and Stanley deal with the implicatures embedded in the antecedents of conditionals in the same way, by exploiting the context-dependence of the similarity relation in terms of which conditionals are standardly analysed.¹⁹

In reply to King and Stanley, let me note, first, that focussing is not necessary to get the desired effect. We can say:

(10) Eating some of the cake is better than eating all of it

¹⁸ In this respect, the 'better-than' construction is similar to conditionals, which also give rise to embedded implicatures.

¹⁹ Gazdar is the first theorist to have attempted to explain away the embedded implicatures of Cohen-conditionals by appealing to the context-sensitive semantics of conditional sentences in the Stalnaker-Lewis framework. See Gazdar 1979: 70. At the very end of the same book, however, Gazdar seems to revert to the view that there are genuine embedded implicatures in examples like (8)-(9).

Even though the word 'some' does not bear focal stress, still a contrast is made (by means of the 'better than' construction itself) between 'some' on the left-hand-side and 'all' on the right-hand-side. Now such a contrast makes sense only if we strengthen 'some' so as to get the upper-bounded reading. The notion of 'making sense' that comes into play here is pragmatic in the fullest possible sense (the 'top-down' sense), hence the appeal to saturation is misguided: the process of strengthening takes place not because it is linguistically mandated, but in order to make sense of the speaker's communicative act. King and Stanley hide this fact by using only examples in which some formal feature of the linguistic signal (viz. the use of focal stress) drives the strengthening process. But this feature is not essential to the example — we can get rid of it, as in (10).

King and Stanley argue that, in this particular case, the strengthening effect can be achieved indirectly by manipulating the similarity relation invoked by the 'better-than' construction, rather than by directly enriching the meaning of 'some'. Granted; but this is, once again, a feature of the example that is accidental and irrelevant to the issue. We will have exactly the same strengthening effect *whichever construction we use*, as long as it involves a contrast between 'some' and 'all'. That will be so, in particular, *even if the construction in question does not invoke a similarity relation between worlds*. For example we may say:

John ate some of his cake but Jim ate all of his

Here the scalar enrichment of 'some' takes place within the first conjunct, in the scope of 'but', yet it cannot be explained away in terms of some process of saturation that independently takes place in interpreting this construction. (There *is* a process of saturation at work in the interpretation of 'but', but I take it to be irrelevant to the matter at stake.)

To conclude, the weakness of the NN position is that it offers no general account of embedded implicatures. It only gives us saturation-based analyses for particular cases — analyses which cannot be generalized because they exploit accidental features of the cases in question.

The next position is the YN ('Yes-No') view, which posits default implicatures but sticks to a classical conception of pragmatics as operating on the output of the

grammar. I take this to be Chierchia's position. This may sound paradoxical, since those who, like Chierchia, posit default implicatures explicitly reject the 'Gricean Root' and hold instead that 'pragmatic computations and grammar driven ones are "interspersed"' (Chierchia 2001: 1). But the pragmatic computations that are said to interfere with semantic composition are not genuine pragmatic processes — the sort of pragmatic process that TCP talks about. What DGSi theorists hold is that there is, *in the grammar*, a mechanism that calculates default implicatures. That mechanism operates locally. But it is not a genuine pragmatic mechanism: it is context-independent and belongs to the linguistic system. As far as genuine pragmatics is concerned, a DGSi theorist may well assume a conservative conception of pragmatics as operating on the output of grammar (where 'grammar' now *includes* the DGSi mechanism).

According to Chierchia, when a default implicature that has been automatically factored in in calculating the truth-conditions of the sentence is felt contextually inappropriate, some backtracking takes place and the implicature is cancelled. This can only happen globally — after the default truth-conditions of the utterance have been figured out. As Chierchia writes, 'cancellation amounts to a simple kind of backtracking'. The speaker computes the default meaning (truth-conditions) of the sentence, increments the context with the sentence thus interpreted, and if that fails, backtracks and gets rid of some default implicature so as to get a more acceptable interpretation. (An alternative procedure is to change the context through accommodation: see below.)

On the YN view, the phenomenon of embedded implicatures is accounted for entirely in terms of defaults. Only default implicatures can be embedded and arise locally; nonce implicatures cannot, for the pragmatic mechanism which yields them is classical, that is, globalist. This view is empirically testable. It will be *prima facie* disproved, if we can find a case in which an implicature which results from a full-fledged pragmatic process is embedded. For example, let us consider cases in which some default implicature which has been suppressed because it occurs in a downward entailing environment is contextually reinstated ('frozen'). In such cases it is clear that the implicature is absent from the output of grammar (since it has been suppressed): its presence in the actual interpretation of the utterance is entirely due to the pragmatic process which overrides the default removal. If that pragmatic process can only be global and post-propositional, as the YN view holds, the

implicature in question cannot be embedded. But it is not difficult to find cases in which a nonce implicature is embedded. Chierchia himself gives the following example:

- (11) It was a two-course meal. But everyone who skipped the first or the second course enjoyed it more, for he wasn't too full to appreciate it.

The disjunction ('skipped the first or the second course') occurs in a downward entailing environment here, so the scalar implicature which normally gives rise to the exclusive reading of 'or' is suppressed. In the default interpretation delivered by the computational system of grammar, therefore, 'or' takes the inclusive interpretation.²⁰ Still, the context is such that only the exclusive interpretation makes sense. As Chierchia says, 'we don't mean to include among the most satisfied customers, people who skipped both courses'. The removed implicature is therefore contextually reinstated. This is a genuine pragmatic process, one that, according to the YN view, cannot take place locally. Yet, as Chierchia notices, the scalar implicature which results in the exclusive reading of the disjunction in (11) cannot be accounted for by the global method. The implicature we would get by the global method would be something like 'Not everyone who skipped the first and the second course enjoyed the meal more!'²¹ This is not at all what we want. Chierchia concludes that 'cases like (11) seem to constitute further evidence against globalism. We seem to be in presence here of an embedded scalar implicature' (p. 8). Insofar as the implicature in question is not produced by the computational system of grammar but by a full-

²⁰ Thus in a standard example like 'Every student who wrote a squib or made a classroom presentation got extra credit', the default interpretation of the disjunction is clearly inclusive.

²¹ The global method consists in first replacing the scalar term by the weakest of its stronger alternatives, and then negating the result. The (only) stronger alternative to 'the first or the second course' is 'the first and the second course', so the 'global' implicature of 'Every one who skipped the first or the second course enjoyed the meal more' ought to be 'Not everyone who skipped the first and the second course enjoyed the meal more'.

fledged pragmatic mechanism, this type of example seems to show that the YN view cannot be right.

Yet Chierchia sticks to the classical view of the relation between grammar and (genuine) pragmatics, so he has to deny that there is an embedded implicature in (11), appearances notwithstanding. To that effect he uses the same strategy as King and Stanley, and argues that we can account for the strengthening effect indirectly, by manipulating the domain of the quantifier through accommodation:

The interpretation of (11) requires a domain of people who don't skip both courses. This yields the same effect that we would obtain by not removing the locally added implicature. (Chierchia 2001: 18)

Chierchia also appeals to that strategy to handle the scalar implicatures that are embedded in the antecedent of conditionals (since the antecedents of conditionals are downward entailing environments in which scalar implicatures are removed by default). Thus he analyses Levinson's example (12) exactly as Gazdar or King and Stanley would:

(12) If John has two cars, the third one parked outside must be somebody else's

'Here too', Chierchia says,

we want to accommodate in the antecedent of (12) an 'and no more' proviso. I.e., we want to restrict our consideration to sets of worlds from which people with more than two cars are excluded. The effect of this accommodation is the same as the computation of an implicature. But if we are right, the mechanism through which this happens is very different from how normally implicatures come about. In (12) the implicature is not added in locally. It is accommodated at some point to avoid a near contradiction. (Chierchia 2001: 18)

The last two positions are NY and YY. NY is the view supported by relevance theorists. Deirdre Wilson has been the first linguist perhaps to notice the importance of embedded implicatures. Examples like (9)-(10), which parallel the Cohen-

conditionals, are originally due to her (Wilson 1975: 151). Insofar as they result from a pragmatic process, embedded implicatures show that such processes can take place locally and affect truth-conditions — a theme central to relevance theory. From the very beginning, relevance theory has made room for primary pragmatic processes such as free enrichment, in terms of which embedded implicatures can easily be accounted for. On the other hand relevance theorists have never been happy with the notion of default pragmatic value. They take the alleged 'defaults' to be not values calculated by the language faculty independent of context, but pragmatic values that are determined (in the usual, context-sensitive way) in the most easily accessible contexts that come to mind when no specific context is otherwise provided. For example, one might argue that the phrase 'cut the grass' has a default reading (in which 'cut' is understood in the specific sense of *mow*) — an enriched reading which can be contextually overridden and which springs to mind when the phrase is produced 'out of the blue'. Relevance theorists would reply that, when the words 'cut the grass' are uttered out of the blue, a stereotypic scenario is evoked which is used as context (default context, we might say) and with respect to which the meaning of 'cut' may undergo free enrichment, thereby yielding the alleged default reading. The mechanism at work in producing that reading is exactly the same as (and is no less context-sensitive than) the mechanism at work in producing a context-specific reading such as *slice the grass into strips* in one of Searle's fancy scenarios.²² According to relevance theorists, the same thing goes for scalar implicatures: the mechanism that is responsible for the alleged 'default implicatures' is exactly the same as the mechanism that is responsible for context-specific, nonce implicatures.

Some work is needed on the part of relevance theorists if they want to substantiate their claim. For there is a significant difference between the default reading of 'cut the grass' and the default reading of an utterance of the form 'If P or Q, then R'. Chierchia cites psychological experiments by Noveck et al. (2001),

²² 'Suppose you and I run a sod farm where we sell strips of grass turf to people who want a lawn in a hurry... Suppose I say to you, "Cut half an acre of grass for this customer"; I might mean not that you should *mow* it, but that you should slice it into strips as you could cut a cake or a loaf of bread' (Searle 1980: 224-25).

establishing that of the two inferences that follow, the first one is overwhelmingly accepted, while the second one is rejected:

(13a)	If P or Q, then R P and Q	<hr/> R
(13b)	If P then Q and R P	<hr/> Q or R

Those experiments seem to confirm that, by default, 'or' is interpreted inclusively in the antecedent of a conditional, but exclusively in a simple statement like the conclusion of (13b). This fact cannot easily be accounted for in terms of stereotypic scenarios. As Chierchia writes:

What is interesting here is that we are dealing with abstract syllogistic frames with letter variables, where the only "real" words are *or* and *if...then*. Hence the relevant effect cannot be imputed to anything like scripts, and lexical or world knowledge of any kind. It must be due to the meaning of the only "real" items that occur in the experimental material. (p. 35)

I turn to the YY view, which has never been explicitly defended in the literature²³ and which I'd like to advertise, in closing this paper. Like YN, YY posits two mechanism, one semantic (the DGSI) and the other pragmatic, but the pragmatic mechanism is taken to be primary rather than secondary: it interacts with the process of semantic composition, rather than operating on its output. To make sense of this position, let us consider the processing model outlined in Recanati 1995 and see how Chierchia's ideas could be accommodated within such a framework.

²³ Stephen Levinson, the leading advocate of default pragmatic values, is also an advocate of Truth-Conditional Pragmatics (he accepts free enrichment, pragmatic intrusion, etc.), so he is a potential client for the YY view.

In that paper I argued that the three contextual processes of disambiguation, saturation, and modulation have the following characteristics in common. In all cases there are several *candidates* for the status of compositional value (the compositional value of an expression-token being what it contributes to the interpretation of the sentence-token where it occurs). In disambiguation the candidates are the distinct meanings of the ambiguous expression. In saturation the candidates are the various things which can be assigned, in context, to the linguistic element in need of saturation; for example, different referents/antecedents can be assigned to the pronoun 'he' in the discourse 'John was arrested by a policeman yesterday; *he* had just stolen a wallet'. In modulation the candidates are (i) the lexical value encoded by the word, and (ii) pragmatically derived values that are freely generated, on an associative basis, by applying various functions to the lexical value. In all cases, the candidates receive activation from multiple sources (including, but not restricted to, linguistic sources), and the candidate that wins (that which actually goes into the interpretation and assumes the status of compositional value) is that which gets the highest activation when the process of activation spreading stabilises.

The important point, for our present purposes, is that the activation which a candidate ultimately receives comes from several sources. One source of activation is linguistic. For example, it is well-known that when an ambiguous word is uttered, its distinct meanings, however contextually inappropriate, are automatically activated. This simply means that part of the activation which a given meaning receives comes from the fact that the word that encodes that meaning has been uttered. Activation also comes from other sources, of course, and if no activation comes from other sources the meaning in question is deactivated and loses the competition. Or consider modulation: the lexical value is automatically activated, and it is through the lexical value that the other candidates are accessed and can themselves get activated. But which value actually gets into the interpretation and becomes the compositional value depends upon the activation level reached by the various candidates when all sources of activation have been taken into account. The lexical value will get into the interpretation only if it receives enough activation from other sources to stay at the top of the activation ranking.

In this light we can see the DGS as doing two things. First, it provides further candidates (the strengthened meanings) over and above the lexical values (the plain meanings). Second, it is a (linguistic) source of activation for the candidates — it

contribute a certain degree of activation to them along the lines of the Chierchia algorithm: the strengthened value of scalars gets a higher degree of activation unless a downward entailing operator is encountered, in which case the ranking is reversed and the plain value gets a higher ranking. Or perhaps, we should construe the strengthened readings as getting a higher activation than the plain readings, and the downward entailing operators as deactivating the strengthened meanings of the scalars in their scope, so as to reverse the initial ranking.²⁴ Whatever the details, the grammar will be a linguistic source of activation for both the plain meaning and the strengthened meaning. But this is only one source of activation, as in the other cases. On the overall picture, grammatical and contextual factors jointly contribute activation to the candidates, in such a way that a scalar term in the scope of a downward entailing operator may be assigned the strengthened meaning if the contextual activation of that meaning is higher than the deactivation resulting from default removal. On this 'common currency model', we don't have to construe the overriding of defaults as an instance of backtracking or anything of the sort. On the YY view, defaults are overridden locally, just as referents and indexical values are assigned locally.²⁵

References

Anscombe, Jean-Claude and Ducrot, Oswald (1978), 'Echelles implicatives, échelles argumentatives et lois de discours', in *Semantikos* 2 (2-3): 43-66.

Anscombe, Jean-Claude and Ducrot, Oswald (1983), *L'Argumentation dans la langue* (Brussels: Mardaga).

Bach, Kent (1994), 'Conversational Implicature', in *Mind and Language* 9: 124-62.

²⁴ This formulation is not quite satisfactory because it does not take the 'flip-flop' effect into account. What downward entailing operators must do is reverse the ranking previously established between plain meaning and strengthened meaning. This amounts to deactivating the strengthened meanings only in certain cases.

²⁵ I am grateful to Gennaro Chierchia, Benjamin Spector and Dan Sperber for discussions or comments.

- Bach, Kent and Harnish, Mike (1979), *Linguistic Communication and Speech Acts* (Cambridge, Mass: MIT Press).
- Bezuidenhout, Anne (2002), 'Generalized Conversational Implicatures and Default Pragmatic Inferences', in J. Kaim Campbell, M. O'Rourke and D. Schier (eds.) *Meaning and Truth: Investigations in Philosophical Semantics* (New York: Seven Bridges Press), 257-83.
- Campbell, Robin (1981), 'Language Acquisition, Psychological Dualism and the Definition of Pragmatics', in H. Parret, M. Sbisà, and J. Verschueren (eds.) *Possibilities and Limitations of Pragmatics* (Amsterdam: Benjamins), 93-103.
- Carston, Robyn (1997), 'Enrichment and Loosening : Complementary Processes in Deriving the Proposition Expressed?', in *Linguistische Berichte* 8: 103-27.
- Carston, Robyn (2002), *Thoughts and Utterances : The Pragmatics of Explicit Communication* (Oxford: Blackwell).
- Carston, Robyn (forthcoming), 'Semantics and Conversational Implicature', in the Proceedings of the Genoa conference *WOC 2002: the Semantics/Pragmatics Distinction*, to be published by CSLI Publications.
- Chierchia, Gennaro (2001), 'Scalar Implicatures, Polarity Phenomena, and the Syntax/Pragmatics Interface' (typescript).
- Cohen, Jonathan (1971), 'Some Remarks on Grice's Views About the Logical Particles of Natural Language', in Y. Bar-Hillel (ed.) *Pragmatics of Natural Languages* (Dordrecht: Reidel), 50-68.
- Cohen, Jonathan (1977), 'Can the Conversationalist Hypothesis Be Defended ?', in *Philosophical Studies* 31: 81-90.
- Cornulier, Benoît de (1982), 'Sur le sens des questions totales et alternatives', in *Langages* 67: 55-109.
- Cornulier, Benoît de (1984), 'Pour l'analyse minimaliste de certaines expressions de quantité', in *Journal of Pragmatics* 8: 661-691.
- Ducrot, Oswald (1969), 'Pré-supposés et sous-entendus', in *Langue Française* 4: 30-43.
- Ducrot, Oswald (1972), *Dire et ne pas dire* (Paris: Hermann).
- Ducrot, Oswald (1973), *La preuve et le dire* (Paris: Mame).
- Fauconnier, Gilles (1976), 'Remarques sur la théorie des phénomènes scalaires', in *Semantikos* 1 (3): 13-36.
- Gazdar, Gerald (1979), *Pragmatics* (London: Academic Press).

- Geurts, Bart (1998), 'Scalars', in P. Ludewig and B. Geurts (eds.), *Lexicalische Semantik aus Kognitiver Sicht* (Tübingen: Gunter Narr), 95-117.
- Green, Mitchell (1998), 'Direct Reference and Implicature', in *Philosophical Studies* 91: 61-90.
- Grice, Paul (1989), *Studies in the way of Words* (Cambridge, Mass: Harvard University Press).
- Horn, Larry (1989), *The Natural History of Negation*, Chicago University Press.
- Jackendoff, Ray (1997), *The Architecture of the Language Faculty* (Cambridge, Mass: MIT Press).
- Karttunen, Lauri (1973), 'Presupposition of Compound Sentences', in *Linguistic Inquiry* 4: 169-193.
- Karttunen, Lauri (1974), 'Presupposition and Linguistic Context', in *Theoretical Linguistics* 1: 181-194.
- King, Jeffrey and Stanley, Jason (forthcoming) 'Semantics, Pragmatics, and the Role of Semantic Content'.
- Landman, Fred (2000), *Events and Plurality* (Dordrecht: Kluwer).
- Levinson, Stephen (1983), *Pragmatics* (Cambridge: Cambridge University Press).
- Levinson, Stephen (2000), *Presumptive Meanings : The Theory of Generalized Conversational Implicature* (Cambridge, Mass: MIT Press).
- Mackie, John (1973), *Truth, Probability and Paradox* (Oxford: Clarendon Press).
- Morgan, Jerry (1978), 'Two Types of Convention in Indirect Speech Acts', in *Syntax and Semantics* 9: 261-80.
- Noveck, Ira *et al.* (2001), 'Linguistic-Pragmatic Factors in Interpreting Disjunctions' (typescript).
- Nunberg, Geoffrey (1995), 'Transfers of Meaning', in *Journal of Semantics* 12: 109-32.
- Nunberg, Geoffrey and Zaenen, Annie (1992), 'Systematic Polysemy in Lexicology and Lexicography', in H. Tommola, K. Varantola, T. Tolonen, and J. Schopp (eds.) *Proceedings of Euralex 2* (University of Tampere), Part II, 387-98.
- Partee, Barbara (1984), 'Compositionality', in F. Landman and F. Veltman (eds.) *Varieties of Formal Semantics* (Dordrecht: Foris), 281-312.
- Recanati, François (1989), 'The Pragmatics of What is Said', reprinted in S. Davis (ed.) *Pragmatics: A Reader* (Oxford: Oxford University Press, 1991), 97-120.

- Recanati, François (1993), *Direct Reference: From Language to Thought* (Oxford: Blackwell).
- Recanati, François (1995), 'The Alleged Priority of Literal Interpretation', in *Cognitive Science* 19: 207-32.
- Recanati, François (2001), 'What is Said', in *Synthese* 128: 75-91.
- Recanati, François (2004), *Literal Meaning* (Cambridge: Cambridge University Press).
- Sag, Ivan (1981), 'Formal Semantics and Extralinguistic Context', in P. Cole (ed.) *Radical Pragmatics* (London: Academic Press), 273-94.
- Searle, John (1975), 'Indirect Speech Acts', in *Syntax and Semantics* 3: 59-82.
- Searle, John (1980), 'The Background of Meaning', in J. Searle, F. Kiefer, and M. Bierwisch (eds.) *Speech Act Theory and Pragmatics* (Dordrecht: Reidel), 221-32.
- Sperber, Dan and Wilson, Deirdre (1986), *Relevance : Communication and Cognition* (Oxford: Blackwell).
- Stalnaker, Robert (1974), 'Pragmatic Presuppositions', in M. Munitz and P. Unger (eds.) *Semantics and Philosophy* (New York: New York University Press), 197-214.
- Stalnaker, Robert (1999), *Context and Content* (Oxford: Oxford University Press).
- Stanley, Jason (2000), 'Context and Logical Form', in *Linguistics and Philosophy* 23: 391-434.
- Walker, Ralph (1975), 'Conversational Implicature', in S. Blackburn (ed.), *Meaning, Reference and Necessity* (Cambridge: Cambridge University Press), 133-81.
- Wilson, Deirdre (1975), *Presupposition and Non-Truth-Conditional Semantics* (London: Academic Press).