June 2007

Revised

June 2007

Drawing the Line Between Meaning and Implicature – and Relating both to Assertion

By

Scott Soames School of Philosophy University of Southern California

Drawing the Line Between Meaning and Implicature -

and Relating Both to Assertion

Paul Grice's theory of Conversational Implicature is, by all accounts, one of the great achievements of the past fifty years -- both of analytic philosophy and of the empirical study of language. Its guiding idea is that constraints on the use of sentences, and information conveyed by utterances of them, arise not only from their conventional meanings (the information they semantically encode) but also from the communicative uses to which they are put. In his view, the overriding goal of most forms of communication is the cooperative exchange of information -- the pursuit of which generates norms for its rational and efficient achievement. Among them are Grice's conversational maxims.

Maxims of Quantity

- 1. Make your conversational contribution as informative as is required (by current conversational purposes). In other words, don't say too little.
- Don't make you conversational contribution more informative than is required.
 Don't say too much.

Maxims of Quality

- 1. Don't say what you believe to be false.
- 2. Don't say that for which you lack adequate evidence.

Maxim of Relevance

Make your conversational contribution relevant to the purpose of the conversation -- i.e. be relevant.

Maxims of Manner

- 1. Avoid obscurity of expression.
- 2. Avoid ambiguity.
- 3. Be brief.
- 4. Be orderly.

Grice uses these maxims to define *conversational implicature*. His picture is this. A person utters a sentence S in a context in which it semantically expresses a certain proposition. In most cases, this proposition is identified with *what the speaker says*. Although this proposition is part of the information conveyed by the utterance, in many cases it does not exhaust this information. In addition, the speaker also *implicates* certain things. Grice's notion of *conversational implicature* is one particularly important kind. The definition goes roughly as follows.ⁱ

Conversational Implicature

A speaker conversationally implicates q by saying p iff (i) the speaker is presumed to be observing the conversational maxims, (ii) the supposition that the speaker believes q is required in order to make his saying p consistent with the presumption that he/she is obeying the maxims, and (iii) the speaker thinks that the hearers can recognize this requirement, and also that they can recognize that he/she knows that they can do so.

Grice illustrates the idea in the following remark.

"The presence of a conversational implicature must be capable of being worked out; for even if it can in fact be intuitively grasped, unless the intuition is replaceable by an argument, the implicature (if present at all) will not count as a conversational implicature. To work out that a particular conversational implicature is present, the hearer will rely on the following data: (1) the conventional meaning of the words used, together with the identity of any references that may be involved [in contemporary parlance, the semantic content of the sentence relative to the context]; (2) the Cooperative Principles and its [conversational] maxims; (3) the context, linguistic or otherwise of the utterance; (4) other items of background knowledge; and (5) the fact (or supposed fact) that all relevant items falling under the previous headings are available to both participants and both participants know or assume this to be the case. A general pattern for the working out of a conversational implicature might be given as follows: "He has said that p; there is no reason to suppose that he is not observing the maxims, or at least the Cooperative Principle; he could not be doing this unless he thought that q; he knows (and knows that I know that he knows) that I can see that the supposition that he thinks that q is required; he has done nothing to stop me thinking that q; he intends me to think, or is a least willing to allow me to think, that q, and so he has implicated that q.ⁿⁱⁱ (31)

Two points are worth noting. First, on the standard picture, a conversational implicature is a piece of information conveyed, over and above what is said, or asserted – which is itself closely related to the meaning of the sentence uttered, and perhaps identical with what would now be called *the semantic content of the sentence in the context*.ⁱⁱⁱ Second, in order for a proposition q to count as conversationally implicated, the conclusion -- that the speaker believes or accepts q, and is inviting his hearers to do the same – must, in principle, be derivable by an argument of the sentence in the conversational meaning of the sentence uttered by an argument of the specified type from information available to speaker-hearers about the meaning of the sentence uttered, the context of utterance, the conversational maxims, and other background information in the context. A proper understanding of these two points is crucial to distinguishing conversational implicature from linguistic meaning, and determining the extent to which this distinction reflects something psychologically real about how speaker-hearers process linguistic information, and interpret linguistic performances.

The relationship between assertion and implicature is illustrated by Grice's own examples. Here are two involving the first maxim of quantity – *Don't say too little*.

(i) A is planning to drive up the coast with B. Both know that A wants to see C, if doing so is feasible.

A: Where does C live?

B: Somewhere in Northern California

B conversationally implicates that he doesn't know where in Northern California C lives – since if he did, he would be violating the first maxim of quantity by not giving the location. Here, the presumption that he is not violating the maxims requires one to conclude that he does not take himself to know where C lives. (Adapted from Grice 1989, 32-33.)

(ii) A is writing a recommendation for a student who is a candidate for a philosophy job, and his letter reads as follows: "Dear Sir, Mr. X's command of English is excellent, and his attendance in class has been regular. Yours, truly." A implicates that his student is no good at philosophy. Why? A knows that a more informative letter is desired. Since he is also in a position to provide the needed information, there must be a reason he hasn't put it in the letter. There is no reason to think he is being uncooperative, since if he were, he wouldn't have written. Given that he is being cooperative, he would surely give us a positive evaluation if he had one to give. (People are more reluctant to state negative evaluations than positive ones, especially in writing.) Thus, his evaluation must be negative, and he simply doesn't want to explicitly say so. Hence, he must think the student is no good. (Adapted from Grice 1989, 33)

In these cases the proposition conversationally implicated is clearly *not* asserted. In (i), the proposition B asserts – that C lives somewhere in Northern California – directly advances the purpose of the interchange, whereas the proposition that B conversationally implicates – that B doesn't know where in Northern

California C lives – is merely collateral information explaining why B chose to say what he did. This fits the classic Gricean pattern. Assertion is conceptually prior to conversational implicature, which adds further information about why the particular assertion, rather than some other, was made. This pattern applies to (ii) as well, even though the reasoning generating the implicature is more convoluted, and the proposition implicated – that the job candidate is no good – is the real point of the writer's remark. Although this may tempt one to identify the implicature as the writer's "real assertion," the temptation should be resisted -- since the whole purpose of using indirect means to convey this information was to avoid having to state it. In cases like this, the speaker or writer *exploits* shared knowledge of the conversational maxims to convey crucial information without asserting it.

This idea is taken a step further by examples involving the first maxim of quality – *Don't say what you believe to be false.*

- (iii) *Irony*: X, with whom A has been on close terms until now, has betrayed a secret to A's rival, a fact known to both A and his hearers. A says *X* is a fine friend. Since this remark is obviously false, either A has violated the first maxim of quality, or he means something else by his remark. Since there is no reason to think that he is opting out of the cooperative principle, he must be intending to convey something else most likely, the opposite. Note, A hasn't asserted the proposition that X is a fine friend, and implicated something further. He hasn't asserted the proposition expressed by his sentence at all. Rather, the presumption that the conversational maxims have been obeyed defeats the normal presumption that the semantic content of the sentence is included in what the speaker says.
- (iv) Metaphor: The speaker says You are the cream in my coffee. Here, the proposition literally expressed by the sentence is false, since you refers to a person. As a result, it is obvious that the speaker can't be trying to convey that proposition. (Quality) Nor is he trying to convey its

negation, the truth of which is too obvious to need conveying. (Quantity) Consequently, he must be trying to convey an idea naturally suggested by the sentence – something along the lines that the person to whom he is speaking is as crucial to him as cream is to his coffee.

((iii) and (iv) adapted from Grice 1989, p. 34.)

In both of these cases, Gricean maxims affect what is asserted. However, it is not clear that they generate implicatures that constitute all, or part, of the speaker's assertion. What the speaker conversationally implicates is that he is not asserting what his sentence literally means (its semantic content in the context), but rather something related to it - which it is up to his audience to discern. In the case of live metaphors – in contrast with the hackneyed example in (iv) – there is often a tantalizing openness in the content of the assertion, lending it a desired depth that challenges listeners in ways that go beyond the banality of standard conversational implicature.^{iv} Cases in which there is an obvious and definite candidate for the proposition asserted – different from the one literally expressed by the sentence uttered - should, I think, be seen as limiting cases of this process. Thus, I don't classify the propositions asserted in (iii) and (iv) as conversational implicatures, even though the explanation of how they come to be asserted relies in part on Gricean maxims. Even if one objects, however, and does view them as asserted implicatures, the way in which implicatures contribute to assertion remains highly restricted. They can defeat the presumption that the semantic content of the sentence uttered is asserted, thereby forcing a reinterpretation that gives rise to non-literal assertion. However, on the classical Gricean story, they do not enter into the assertive contents of normal, literal uses of language.^v

This is related to the second point noted earlier – namely the requirement that conversational implicatures be derivable by speaker-hearers, using the standard Gricean form of argument, from their knowledge of the meaning of the sentence uttered, and their recognition of what, literally interpreted, the utterance can, or would, be taken to assert (plus their awareness of the conversational maxims and the

shared assumptions in the context). For Grice, the purpose of the derivability requirement is (i) to distinguish conversational implicatures from other information, either present in the speech situation (e.g. that someone is speaking) or conventionally encoded by the sentence uttered, (ii) to establish conversational implicature as a normative notion, rather than a purely descriptive one about the psychological processes at work in the real-time extraction of information from utterances, and (iii) to ground this normative notion in knowledge that speaker-hearers actually have. Although (i) and (ii) are well-motivated and, in my opinion, quite correct, there is a legitimate worry that the knowledge required by (iii) may exceed what competent speaker-hearers can reasonably be expected to have.

The standard Gricean story tacitly assumes that meaning and assertion are basically transparent – in the sense of being reliably identifiable independent of whatever further, non-semantic, nonassertive content may be carried by an utterance. The Gricean argument grounding every conversational implicature *starts* with the presumed recognition by speaker-hearers of meaning and assertive content – *what is said*. From there it proceeds via the cooperative principle and conversational maxims, plus (in some cases) special features of the context, to the conclusion that the speaker must be relying upon, and intending to communicate, some further proposition, acceptance of which is required by adherence to the maxims. Although it would be unrealistic to suppose that speaker-hearers actually construct derivations of this sort every time an utterance conversationally implicates something, Grice, quite rightly, doesn't require this. What he does require is that every implicature be, in principle, derivable from information that speaker-hearers really possess. The worry is that they may not always possess the information about meaning needed to get the required derivations started.

Grice expresses a version of this worry in connection with one of his most successful, and theoretically informative examples. In logic, a disjunction is true iff either one or both of the disjuncts are true; it is false otherwise. However, philosophers have sometimes wondered whether *or* has the

same meaning in ordinary language as it does in logic.^{vi} One reason for this comes from the fact that nearly always when we assertively utter [A or B], we do so *not* because we already know that A is true, or because we already know that B is, but rather because we have reason to believe that it is unlikely, perhaps impossible, that both A and B will fail to be true. Though we don't know whether A is true, and we don't know whether B is true, we have grounds for thinking that one will be true, should the other turn out not to be. So, when we assert a disjunctive proposition, we generally have non-truth-functional grounds for doing so – grounds which are not grounds for asserting the proposition expressed by either disjunct. This is such a pervasive feature of our ordinary talk that it might seem to be a special feature of the meaning of *or* in ordinary language that is not part of its meaning in formal logic.

Grice argues that this feature of the use of *or* is matter of conversational implicature, not meaning. Suppose one assertively utters [A or B] in a conversation in which one is abiding by the conversational maxims. Suppose further that *or* is the logician's *or*, and so [A or B] has the standard truth conditions given in formal logic. Then, the proposition asserted is weaker -- less informative -- than the proposition expressed by A -- as is shown by the fact that it is entailed by that proposition, but not vice versa. The same is true of B. The first maxim of quantity directs one to make the strongest, most informative statement one can, provided that it would be relevant, and provided also that doing so would not violate other maxims -- including the maxim not to assert that for which one lacks adequate evidence. Since one has asserted the weaker statement – expressed by the disjunction -- rather than the statement expressed by either disjunct, and since, presumably, if the disjunction is relevant to the conversation, then each disjunct is too, the presumption that one is obeying the conversational maxims requires one's hearers to conclude that one lacks adequate evidence to assert the statement expressed by

either disjunct. But since one must have adequate evidence to assert the statement expressed by [A or B], one must have non-truth-functional grounds for it.

So, someone who asserts the proposition expressed by [A or B] typically conversationally implicates that he/she doesn't know, or have adequate grounds to assert, the proposition expressed by A, or the proposition expressed by B, but that he/she *does* have adequate grounds for thinking that both can't be false. This is the feature of ordinary use that led to the suspicion that *or* has a sense in ordinary language different from its sense in logic. What Grice has shown is that this suspicion is unfounded. Even if we start with the logical *or*, we will end up concluding that assertive utterances employing it conversationally implicate precisely what is implicated by utterances of ordinary disjunctions. Thus, we have no reason to posit a special sense of *or* in ordinary language, distinct from its sense in logic.

Although this argument is decisive, Grice uses it to raise a larger, more troubling issue. He imagines someone remaining unconvinced, and maintaining that it is intuitively clear that the sense of *or* in ordinary language differs from its sense in logic. Though the particular complaint is unreasonable, the larger issue is not. How, if Grice is right, could there ever be controversy about what information is part of meaning, and what is conversationally implicated? The worry is that the Gricean model would not seem to allow for serious controversy of this sort. It is a basic presupposition of the model that the generation of conversational implicatures requires competent speaker-hearers to have a clear, reliable, antecedent grasp of the meaning of the sentence uttered, and of what is said by an utterance of it. But then, it would seem, we should be able to rely on that grasp to determine whether a given piece of information conveyed by the utterance is, or is not, conversationally implicated. If it is part of what is said, this should be recognized by competent speakers, and if it is conversationally implicated, there should be a canonical derivation demonstrating it to be (as in the case of *or* just given). How, then, could matters of this sort ever be controversial?

Grice was both aware that they could be and worried that he hadn't explained how this could be so. Here is how he expressed the worry.

"We must of course give due (but not undue) weight to intuitions about the existence or nonexistence of putative senses of a word (how could we do without them?). Indeed if the scheme which I have been proposing is even proceeding in the right direction, at least some reliance must be placed on such intuitions. For in order that a nonconventional implicature should be present in a given case, my account requires that a speaker shall be able to utilize the conventional meaning of a sentence. If nonconventional implicature is built on what is said, if what is said is closely related to the conventional force of the words used, and if the presence of the implicature depends on the intentions of the speaker, or at least on his assumptions, with regard to the possibility of the implicature being worked out, then it would appear that the speaker must (in some sense or other of the word *know*) know what is the conventional force of the words which he is using. This indeed seems to lead to a sort of paradox: If we, as speakers, have the requisite knowledge of the conventional meanings of sentences we employ to implicate, when uttering them, something the implication of which depends on the conventional meaning in question, how can we, as theorists, have difficulty with respect to just those cases in deciding where conventional meaning ends and implicature begins? If it is true, for example, that one who says *that A or B* implicates the existence of non-truth-functional grounds for A or B, how can there be any doubt about whether the word "or" has a strong or weak sense?" (Grice 1989, 49)

Although I don't think there can be any serious doubt about *or*, I do think that Grice's paradox is genuine, and points to a potential weakness in his account. The weakness is that of too closely identifying what is said or asserted with the meaning (or semantic content) of the sentence uttered, and,

in so doing, failing to recognize that the latter is too theory-laden to be the object of reliable, systematic knowledge on the part of competent speakers. One consequence of this weakness has been to hide the role of the conversational maxims in filling the gap between meaning and assertion. As paradoxical as it may sound from a traditional Gricean point of view, I will argue that there are clear cases of the literal use of language in which what is asserted is jointly determined by the conversational maxims plus the semantic content of the sentence uttered, via Grice-like derivations, despite the fact that these derivations are *not* reconstructable by competent speakers, because the semantic contents that are their starting points are psychologically unavailable to them.

In the next section, I will present a case involving numerical quantifiers that raises substantial theoretical questions about the relationship between meaning, assertion, and conversational implicature – questions which can't be answered by direct appeal to the knowledge of competent speakers. In the sections that follow, I will try to answer these questions, and use the answers to indicate how to draw the line between meaning and implicature. I will close by sketching a broad theoretical framework relating semantic and pragmatic theories, and distinguishing both from psychological theories of the linguistic competence and performance of speaker-hearers.

Implicatures and What is Said: How many?

I begin with sentence (1), assertive utterances of which normally carry the information that the speaker has exactly two children.

1. I have two children.

How much of this information is part of meaning, and how much is due to implicature? Does the sentence literally mean that one has exactly two children, or does it mean that one has at least two children, and conversationally implicate that one has at most two children as well -- or is there a third alternative? The answer to these questions can't be decided by untutored intuition.

Matters become clearer when one looks at related sentences, such as (2).

2. I have two beers in the fridge.

A systematic semantics should treat these sentences, and all other instances of (3), in the same way.

3. I have two Fs.

Either all these instances should mean that one has exactly two of Fs, or none should. Looked at in this way, it seems evident that they don't mean that one has exactly two.

Consider a case in which someone interested in having a quick beer with me asks *Do you have two beers in the fridge*? In the imagined case, it is clear that my guest wants to know whether I have at least two. Hence, when I answer, *Yes, I have two beers in the fridge*, and he goes to the fridge and sees I have six, it won't occur to him that I have said anything false, or misled him in any way. Since he could reasonably think this if my use of the sentence literally meant, and I actually asserted, that I have exactly two beers in the fridge, there is reason to think that (2) doesn't mean this, and that, in the case imagined, I didn't assert or even implicate it.

In different circumstances, the information conveyed would be richer. If a guest, wondering whether there is enough for a party, asks, *How many beers do you have in the fridge?*, I would assume he wanted to know the precise number. When I answer, *I have two beers in the fridge*, I expect him to understand me as saying that I have exactly two. Since the meaning of sentence (2) doesn't change from one conversation to the next, and since no semantic ambiguity seems to be involved, the difference in the cases appears to be non-semantic. In both cases, my utterance carries the information that I have at least two beers in the fridge. When and only when it is conversationally relevant precisely how many beers I have, the utterance carries the additional information that I have at most, and hence exactly, two. Presumably a similar story could be told about (1), and other instances of (3).

Why is this so? One potential explanation is that (3) means (4a)

4a. I have at least two Fs,

and the extra information – that one has at most two Fs -- is a conversational implicature generated by the maxims of quantity and relevance, when the stronger claim

4b. I have exactly two Fs

is conversationally required. However, this is doubtful.

First, utterances of (4a) typically have different conversational implicatures from those of (3). Moreover, the implicatures carried by utterances of (4a) are at variance with those required by the proposed analysis of (3). Standardly, if I assertively utter (4a) in a context in which it is relevant, not just that I have a minimum of two Fs, but also how many I have, my utterance won't conversationally implicate that I have at most, and hence exactly, two. Rather, it will conversationally implicate that (for all I know) I might have more than two Fs. Thus, if I were to say *I have at least two children* in reply to the question *How many children do you have?*, I would typically implicate that I was uncertain regarding my possible parentage of other children. The Gricean explanation of this is transparent. Since the first maxim of quantity requires me to make the strongest, relevant statement I can, I will make the weaker claim expressed by (4a), instead of the stronger (relevant) claim expressed by (4b), only when I lack sufficient evidence for (4b). Instances of (3) don't exhibit this behavior. Thus, far from explaining the information conveyed by utterances of (3), the proposal that (3) means (4s) seems to be an obstacle to such an explanation.

Perhaps, however, this isn't conclusive.^{vii} Let's assume that the semantic contents of (5a) and (5b) are the same (or trivially equivalent).

5a. I have at least one F

5b. I have an F

Still, utterances of them can carry different implicatures. If someone asks, *How many children do you have?*, and I answer, *I have at least one child*, then, as before, I typically will implicate that I might have more. Since this implicature wouldn't be carried by a corresponding utterance of *I have a child*, it is possible for two sentences to have equivalent semantic contents, even though utterances of them differ regarding some implicatures of this sort. Thus, it might be argued, the fact that (3') and (5a) generate different implicatures doesn't show that they differ in meaning either.

(3') I have one F

Similarly for (3) and (4a).

There is, however, a different reason for distinguishing (3') from (5a) and (5b) -- and, by extension, (3) from (4a). If in answer to, *How many Fs do you have?*, I were to reply, *I have an F*, one might naturally wonder whether I had answered the question at all. One would recognize my commitment to having at least one F, while wondering why I left it indefinite whether I have more. At best, one might speculate that -- since I knew that more information was wanted – I was, perhaps, suggesting that I have one and only one F. But one wouldn't judge me to have asserted that I had exactly one F. ^{viii} This contrasts with (3'). If, having been asked *How many Fs do you have?*, I were to reply by assertively uttering, *I have one F*, it would be obvious that I had asserted that I have exactly one F. This difference between (3'), on the one hand, and (5a)/(5b), on the other, remains unexplained -- and, as far as I can tell, unexplainable -- on the assumption that all three mean the same thing. But if (3') doesn't mean (5a), then (3) doesn't mean (4a).

These observations provide tentative support for two conclusions.

C1. Although utterances of sentences of the form (3) typically convey the information expressed by corresponding sentences of the form (4a), and sometimes convey that

expressed by sentences of the form (4b), instances of (3) don't unambiguously mean the same as corresponding instances of (4a) or (4b).

C2. In contexts in which it is conversationally relevant how many Fs one has, assertive utterances of (3) result in assertions of propositions expressed by (4b). Although it is plausible to suppose that the Gricean injunction to make the strongest relevant statement one can plays some role in the assertions made in these contexts, something other than the standard model of conversational implicature is at work.^{ix}

The Case Against Ambiguity

At this point, a simple thought is likely to occur. Perhaps instances of (3) are semantically ambiguous, having one reading in common with instances of (4a), and one in common with instances of (4b). The Gricean injunction to make the strongest, relevant statement one can might then be seen as selecting the stronger reading in cases in which it is conversationally demanded. This view is capable of accommodating all the cases so far considered.

It does, however, raise certain questions. One of these involves the systematicity of the required ambiguity. In our examples, numerals occur as adjectival constituents of quantifier phrases, n Fs, which lack further indication of their precise quantificational force.^x The proposal is that these bare quantifier phrases are ambiguous between two different ways – *at least n Fs* and *exactly n Fs* -- of filling them out. Since there are infinitely many such phrases -- *one F, two Fs, three Fs*, etc. – the proposal requires an infinite number of parallel ambiguities. There is a question about how these are to be understood. If the ambiguities are seen as lexical, there is a danger of losing obvious generalizations. Whereas the different readings of 'bank' are idiosyncratic facts about it (not shared by related words such as 'depository', 'savings and loan,' and the like), the different alleged meanings of *two Fs* are not idiosyncratic in this way. Thus, if an ambiguity is proposed, some explanation of its systematicity will be needed. In addition,

lexical ambiguities typically don't extend across languages. The fact that 'bank' is ambiguous in English carries no expectation that words in other languages with one of its readings will share its other readings as well. Thus, if it turns out that sentences of other languages are routinely used in just the way that instances of (3) are, then this further dimension of systematicity will also have to be explained. Finally, even in English, the proposed ambiguity of instances of (3) can't be restricted to readings shared by instances of (4a) and (4b).

This is shown by the sentences in (6).

6a. Matriculated students are allowed to take five courses.

b. Matriculated students are allowed to take up to five courses.

c. Matriculated students are allowed to take up to, but no more than, five courses.

In many contexts, someone assertively uttering (6a) would be understood as asserting the proposition expressed by (6b) or (6c). The difference between these two is that whereas both propositions stipulate that matriculated students may take any number of courses from one to five, only the latter forecloses taking more. This difference is all but invisible in contexts in which it is tacitly understood that anything not said to be permitted is forbidden. However, in other contexts, the two can be teased apart. For example, although the use of (6b) in D is fine, the use (6c) isn't.

D. I don't know all the regulations, but I do know that matriculated students are allowed to take up to five courses – i.e. they can take any number of courses one through five.
 Whether or not they are allowed to take more, I can't say.

Since in many contexts (6a) can be substituted for (6b) in D without substantial effect, it would seem that the ambiguity theory must posit at least a three-way ambiguity -- with n Fs, being ambiguous between *at least n Fs, exactly n Fs,* and *up to n Fs.*

This point is underscored by comparing (6a,b) with (6d,e).

6d. Matriculated students are allowed to take at least five courses.

6e. Matriculated students are allowed to take exactly five courses.

I assume that *A* is allowed to take at least / exactly / up to five *Fs* is equivalent to *It* is allowed that A take at least / exactly / up to five *Fs*. Since these sentences include numerically quantified phrases in clauses to which the operator *it* is allowed that is attached, scope ambiguity is possible. On the wide-scope readings of the quantifiers, (6d)/(6e) are equivalent to (6dw)/(6ew).

6dw. At least five courses x are such that matriculated students are allowed to take x.

6ew. Exactly five courses x are such that matriculated students are allowed to take x.

These sentences stipulate that the courses that admit matriculated students are at least, or exactly, five in number. They, therefore, entail similar wide-scope readings of (6f).

6f. Matriculated students are allowed to take at least one / two / three / four courses.

However, they do not count the number of courses that may appear on the schedules of matriculated students. Thus, we may put these readings aside.

By contrast, the narrow-scope readings, (6dn)/(6en) of (6d)/(6e) do count the number of courses that may appear on student schedules.

- 6dn. It is allowed that: matriculated students take (a total of) five or more courses i.e. scenarios that make *matriculated students take (a total of) at least five courses* true are permitted.
- 6en. It is allowed that: matriculated students take (a total of) exactly five courses i.e. scenarios that make *matriculated students take (a total of) exactly five courses* true are permitted.

Whereas these readings don't entail the corresponding narrow-scope readings of (6f), the reading of (6b) we are looking for does -- as does the proposition asserted by utterances of (6a) in contexts in which (6a) can be substituted for (6b). This suggests that the ambiguity theory must posit another meaning. Certainly, this use of (6a) can't be explained by taking it to mean (6dn) or (6en), while generating the

additional information carried by it as conversational implicatures. To suppose otherwise, in the case of (6dn), would be to suppose that utterances of (6a) carry both the semantically encoded information that student schedules with five or more courses are allowed, and the conversationally implicated information that schedules with less than five are also allowed. Surely, however, the message conveyed by these utterances is not that promiscuous. The corresponding supposition in the case of (6en) fails to explain (i) why utterances of (6a) should conversationally implicate that schedules with less than five are allowed, when utterances of (6e) do not, and (ii) why any such implicature, even if it could be generated, should be part of what is asserted by an utterance of (6a) – as opposed to being merely suggested by it.

Can we conclude that the ambiguity theory requires (at least) a three-way ambiguity -- with n Fs, being ambiguous between *at least n Fs, exactly n Fs,* and *up to n Fs*? Eventually, yes. But before we do, there is another reading -- illustrated by (6g) -- to consider.^{xi}

6g. Matriculated students are allowed to take at most five courses.

The relevant reading of (6g) is one that specifies five as the most courses that students are allowed to take. This is not the wide-scope reading, since that reading doesn't count the number of courses on allowable schedules. Nor is it the narrow-scope reading, which is equivalent to (6gn).

6gn. It is allowed that: matriculated students take (a total of) n courses, for some $n \le five$.

Rather, it is a reading equivalent to (6g*).

6g* Five is the largest number n such that it is allowed that: matriculated students take (a total of) n courses.

A corresponding reading of (6e) is equivalent to (6e*), which differs from (6en) in ruling out course schedules with more, or fewer, than five courses.

6e*. Five is the only number n such that it is allowed that: matriculated students take (a total of) n courses.

The reading of (6b) we are after is equivalent to $(6b^*)$.

 $6b^*$. For any number n up to five, it is allowed that: matriculated students take (a total of) n courses. This reading is distinct from (6g*) in two respects: (i) unlike (6g*), (6b*) doesn't entail that schedules with more than five courses are disallowed, (ii) also unlike (6g*), (6b*) does entail that course schedules with less than five courses are allowed. More generally, the truth conditions of (6b*) differ from those of the any of the readings of (6d), (6e), and (6g).^{xii}

The crucial point in our argument can now be made. Utterances of (6a) can be used to assert the proposition expressed by (6b*). We have already seen that this can't be explained by assigning (6a) the readings (6dn) or (6en), while appealing to conversational implicatures to fill in the extra needed content. The same is obviously true for the *readings of (6d) and (6e). Hence, the ambiguity theory must posit (at least) a three-way ambiguity – with *n Fs* being ambiguous between *at least n Fs*, *exactly n Fs*, and *up to n Fs*. Other constructions that can be used to make a similar point are given in (7).^{xiii}

7a. The sedan was designed to carry five passengers.

b. The car seats five adults.

c. John is capable of eating five hamburgers at one sitting.

d. The navy can withstand the loss of five carriers without jeopardizing its mission.

At this point, the ambiguity theory should be looking less attractive. Having found a minimum of three required meanings, we have no guarantee that more won't be needed. This is potentially ominous. The greater the number, variety and systematicity of the different contributions made by bare numerical quantifiers to the propositions asserted by utterances of sentences containing them, the less plausible it is to suppose that what we are dealing with is ambiguity.

This point is given additional force by the sentences in (8).

8a. Matriculated students are required to take three courses, and allowed to take five courses.

b. Matriculated students are required to take at least three courses and allowed to take up to five courses.

As before, utterances of the (a) sentence, containing bare numerical quantifiers, may result in assertions of the proposition expressed by the lexically more explicit (b) sentence. This proposition specifies taking three, four, or five courses as allowable, while prohibiting taking only one or two -- the prohibition coming from the first conjunct, and the remaining options from the second. How is this possible? As we have seen, in many discourses uttering (6b) results in the assertion of a proposition which states that matriculated students are allowed to take any number of courses from one to five. But if, under the ambiguity theory, the second conjunct of (8b) is given this interpretation, then it will stipulate one, and two, course loads for matriculated students as allowable – thereby rendering (8b) contradictory. Since (8b) isn't contradictory, something has gone wrong.

In order to diagnose what, we need to look at how permission statements interact with other statements of permission or requirement. When one says that (up to) n courses are permitted, what one says is compatible with some schedules in that range being disallowed on independent grounds, even though they are numerically correct. This is illustrated by (8c).

8c. Matriculated students are allowed to take up to five courses. For students in their first two years, at least one of these courses must be a general education requirement, while for seniors in their final semester, one must be a capstone course in the student's major.

This discourse tells us that course schedules consisting of up to five courses will not be disallowed for having too many, or too few, courses, while indicating that they may still be disallowed if their composition violates other requirements. From this, we conclude that the truth conditions of (8d) are, roughly, those of (8e).

8d. A is allowed to take at least n Fs / exactly n Fs / up to n Fs

8e. *Some scenarios* in which A takes (a total of) at least n Fs / exactly n Fs / up to n Fs are allowed. When (6b*) is understood in this way, it states that for any number n between one and five (inclusive), some schedules with n courses are permitted. Thus, when the initial sentence of (8c) carries this interpretation, the discourse as a whole is perfectly coherent.

However, (8b) is still problematic. Since its first conjunct disallows all schedules with fewer than three courses, it will still be contradictory, if its second conjunct entails that some schedules with less than three are permitted. In principle, the contradiction could be avoided by assigning the second conjunct the narrow-scope interpretation (6bn), which requires only that some schedules with five or fewer courses are allowed.

6bn. Some scenarios in which matriculated students take (a total of) up to five courses are allowed. However, on this reading, the proposition expressed by (8b) is rather weak – since it entails neither that some four-course schedules are allowed, nor that some five-course schedules are. There may be special contexts in which (8b) is used in this way. However, the statement it would normally be used to make does have these consequences, without being contradictory. What it tells us is that three, four, and five course schedules are allowed, while one and two course schedules are not.

Although the understanding we are looking for is simple enough, it requires changing our conception of what *A takes up to n Fs* means. Instead of meaning that the number of Fs that A takes is anywhere from 1 to n (inclusive), it means (roughly) that the number of Fs that A takes is anywhere from a contextually given k to n. In our example, uttering the first conjunct of (8b) changes the default setting of k from 1 to 3^{xiv} In other contexts, k may take on different values. Thus, *up to n Fs* is indexical, or context sensitive. Since the ambiguity theory must posit this indexical meaning as one of the readings of *n Fs*, it must treat these phrases as both ambiguous and context sensitive. However, once we have context sensitivity, ambiguity may be eliminable.

The Case for Semantic Incompleteness and the Contribution of Conversational Maxims' to Assertion

What follows is a different way of looking at the examples we have been concerned with, plus a set of related phenomena – from the larger perspective of an emerging theory of the relationship between semantics and pragmatics, and the relationship between the two. Although both the new theory, and its application to particular cases are somewhat open-ended, each can, I think, be used to shed light on the other. With that in mind, let's return to instances of (9).

9. I have n Fs.

As in our other examples, the numerical quantifier phrase in (9) can be syntactically expanded in a variety of ways: e.g., at least n, exactly n, up to n, at most n, up to n but more than m. Call the contents of these completions possible pragmatic enrichments of the semantic content of the bare numerical quantifier. On the new picture, the semantic content of (9) is free of any such completion, and so is nonspecific about how n Fs should be modified.^{xv} The general constraint governing assertion is that what is asserted by a normal, literal use of a sentence S must be an obvious and relevant pragmatic enrichment of the semantic content of S. xvi When this content is a complete, truth-evaluable proposition, we let the proposition itself count as one of its own possible enrichments. In cases like (9), in which the semantic content is not truth-evaluable, nonvacuous pragmatic enrichment is required in order for a proposition to be asserted. Either way, the semantic content of a sentence can be thought of as a set of conditions that constrains what is asserted by utterances of it. When these constraints fall short of determining a complete, truth-evaluable proposition, pragmatic enrichment is required. When they do determine a complete proposition, enrichment is optional. x^{xvii} If the option is taken, the proposition semantically expressed by the sentence is itself asserted only when it is an obvious and relevant consequence of the enriched proposition that it is the speaker's primary intention to assert.^{xviii}

Pragmatic enrichments of the semantic content of the sentence uttered are influenced by many factors, including the direction and purpose of the conversation, remarks previously made, salient information about the context of utterance, background assumptions shared by speakers and hearers, and Gricean maxims. In addition to generating standard conversational implicatures, over and above what is asserted, the maxims help determine what an utterance asserts by narrowing the class of possible enrichments to those that most effectively advance the conversation. When several enrichments are otherwise feasible, the maxims dictate that one select the strongest, most informative, and relevant propositions among them for which one has adequate evidence. In this way, the maxims of quantity, quality, and relevance play a role in determining what is asserted, and thereby contribute to the truth conditions of utterances.

With this in mind, we return to our earlier examples. In discussing instances of (3) - I have two Fs - I distinguished cases in which the sentence is uttered in response to an interest in how many Fs one has – How many children do you have? -- from cases in which the precise number is unimportant, so long as one has two or more – Do you have two beers in the fridge? The reason one asserts that one has exactly two Fs in one case, while asserting only that one has at least two Fs in the other, is explained by the injunction to make the strongest, relevant statement one can. Since in the former case, the precise number of Fs *is* relevant, one would violate the first maxim of quantity – Don't say too little – if one asserted only that one had at least two Fs. Hence, the stronger enrichment is selected. Since in the latter case the only point at issue is whether one has two or more Fs, one would be giving more information than is needed, if one asserted that one had exactly two Fs. Hence, the conceptually stronger – but conversationally irrelevant – pragmatic enrichment is rejected, and the weaker enrichment is selected in its place.

Next consider (6a) -- *Matriculated students are allowed to take five courses* – understood as a comment on how many, or few, courses students can take. Here, the possible enrichment *at least five* (on

either the narrow-scope or *readings) would result in an assertion that was silent on the permissibility of schedules that would normally be of most interest – namely, those involving one, two, three, or four courses. It is, therefore, not surprising that we don't naturally hear utterances of (6a) in this way. The possible enrichment *exactly five* isn't much better. To learn that matriculated students are allowed to take exactly five courses is to learn that schedules with five courses are allowable – with either no indication of whether schedules with any other number are allowed (the narrow-scope reading), or the rather strange implication that schedules with either more, or fewer, courses are disallowed (the *interpretation). Although there may be possible situations in which the former, very weak assertion, or the latter, unusually strong one, is what is wanted, they are well outside the norm, and do not readily spring to mind.

By contrast, when the enrichment is *up to five*, the resulting proposition – understood as allowing student schedules of one, two, three, four, and five courses -- is strong, natural, and relevantly informative. Thus, adherence to the Gricean maxims leads us to expect what we, in fact, find – namely that utterances of (6a) are naturally understood as making assertions equivalent to (6b*). For the same reason, it is no mystery why, when a truncated version of (6a) occurs as the second conjunct of (8a) -- *Matriculated students are required to take three courses, and allowed to take five courses* – the contribution of that conjunct to the proposition asserted is not the one that would often be asserted when (6a) is uttered on its own. Since if it were, the proposition asserted by uttering the conjunction would violate the maxim against saying what one knows to be false, a different pragmatic enrichment is selected – which, in conformity with the Gricean maxims, yields the strongest, most informative and relevant statement for which the speaker has sufficient evidence. Thus, the theory accommodates the cases we have examined.^{xix}

It may also be possible to use it to accommodate a different class of examples often cited as instances of "loose talk." To take just one case, my home in Southern California is exactly 397 miles from San Francisco by car (taking the shortest route). If, when preparing for a trip, my wife asks *How*

far is San Francisco?, and I answer by assertively uttering (10a), my response would, by normal standards, be judged to be correct.

10a. San Francisco is 400 miles from here.

This will follow, if what I assert is truth expressed by (10b).

10b. San Francisco is approximately 400 miles from here.

On the model we are considering, the semantic content of (10a) is incomplete. Thus, pragmatic enrichment of the bare numerical quantifier is needed to arrive at the proposition asserted. The possible enrichment *exactly 400* seems too precise for the purposes at hand, *up to 400* is too promiscuous, and *at least 400* is too forgiving of distances over 400, and too unforgiving of distances slightly less. Thus, it is natural to take the expected enrichment to be *approximately 400*. In this way, I am characterized as saying something true and nothing false.^{xx}

This result fits the ordinary judgments of speakers. By contrast, the view that I have said (asserted) that San Francisco is exactly (or at least) 400 miles away, and also (somehow) suggested or implicated that it is only approximately that far away seems forced. A "loose-talk" theorist who holds that the semantic content of my remark is the proposition expressed by (10c) or (10d), faces the challenge of explaining (i) which of those two is the real semantic content, (ii) why that content is not heard by speakers as being asserted, and (iii) how the alleged implicature is generated.

10c. San Francisco is exactly 400 miles from.

10d. San Francisco is at least 400 miles from here.

Although some story along these lines could be attempted, there seems to be no need, since a smoother, more natural account is provided by the model used to account for (1-9).^{xxi}

In all these cases, meaning, or semantic content, is incomplete, with conversational maxims playing an important role in selecting contextual completions. In this way, the maxims contribute to what is asserted, and, thereby, to the truth-conditions of utterances. As I have stressed, these cases contrast with classical conversational implicatures, in which one says or asserts one thing, and, as a consequence, implicates something else that is not asserted. Grice's professor, writing a letter of recommendation for a job candidate in philosophy is a case in point. He writes, "Dear Sir, Mr. X's command of English is excellent, and his attendance in class has been regular.," and signs off. In so doing, he implicates, but does not assert, that the student is no good at philosophy. Here, no completion of the semantic content of the sentence is needed to arrive at a complete proposition, and no elaboration of that content by expansion of its constituents could possibly yield the implicature. As a result, the implicated proposition is not asserted, but merely suggested by what is.^{xxii}

Finally, the ideas behind this analysis of the interplay between meaning, assertion, and conversational maxims in cases involving numerical quantifiers have applications beyond those considered here. On this view, the meaning of an expression constrains its contributions to the assertions made by normal, literal utterances of sentences containing it, without always fully determining those contributions. This is what we find with demonstratives like *he*, *she*, *that*, *then*, *there*, *we*, and *now*. Roughly put, the referent of *he* is constrained to be male, the referent of *she* to be female, the referent of *that* to be a non-animate thing , the referent of *then* to be a time, the referent of *there* to be a place, the referent of *we* to be a group that contains the speaker, and the referent of *now* to be a stretch of time including the present moment. To know the meanings of these terms is to know these constraints, and to know that when one uses them in simple sentences one asserts singular propositions about their referents. Everything else used to determine those referents, and to fill out the assertions made by utterances of sentences containing them, is pragmatic enrichment of a sort similar to that found with bare numerical quantifiers.

Thus, it is to be expected that conversational maxims often play a role in determining the propositions asserted by utterances of sentences like (11).

11. He is manic.

Here, the semantics of *he* tells us to select a contextually salient male, without telling us which one. Sometimes overt demonstrations, consisting of pointings, nods, or other gestures, are needed. In other cases, the fact that a certain man m is uniquely relevant to the conversation can be used to fill out the content of the speaker's remark, without any demonstration. If the conversation has been about m, and his psychological state is at issue, the speaker may often rely on the presumption that she is obeying the maxim of relevance to provide the warrant needed for her audience to identify the proposition that m is manic as one she is asserting. A similar result can be achieved, if m is present, and has captured everyone's attention by cavorting in an extravagant manner. In such cases, an assertive utterance of (11), in the absence of any demonstration, will be taken to be a statement about m, for the simple reason that only by taking it to be so will the speaker's remark be relevant to what is going on.

Semantics, Pragmatics, and Psychology: The Nature of Our Theories

It is time to bring these results together, and to illustrate some larger lessons. We started by considering instances of *I have two Fs*. Speakers who understand these sentences are pretty reliable judges about what they, themselves, mean and assert, as well as what others mean and assert by particular utterances. However, they don't have reliable intuitions about whether these sentences mean the same as instances of *I have at least two Fs*, *I have exactly two Fs*, or *I have up to two Fs*. They don't know whether *I have two Fs* means the same as one of these three, is ambiguous between them, or whether it means something else. They certainly don't know what, if the argument in the previous section is correct, the truth of the matter is — namely, that the semantic content of *I have two Fs* in any given context constrains what it is used to assert, without fully determining a truth-evaluable proposition.

Like many correct statements about linguistic meaning, this may well be a theoretical truth to which the practical ability to speak and understand the language gives one no privileged access.

For the same reason, speakers are often unable to separate out, from the total information carried by an utterance, the information contributed by the meaning of the sentence uttered, and distinguish it from that which is added by taking conversational maxims into account. It's obvious that they don't consciously go through Gricean reasoning in all relevant cases. Do they reason it out unconsciously – starting with a psychologically real representation of linguistic meaning, and producing an unconscious Gricean derivation that yields conclusions about what is asserted, and otherwise implicated? Although it's possible they do, I know of no evidence to support such extravagant speculation. Fortunately, there is a natural conception of semantics and pragmatics that doesn't require it.

According to this conception, theories in these domains provide rational reconstructions of the ability of speaker-hearers to identify what is asserted and conveyed by an utterance. The process begins when a sentence is uttered in a context in which many things are known or assumed. As a result of this, the hearer identifies certain things as having been asserted, and other things as having been implied, or suggested. Although the causal, psychological process by which this identification takes place is revealed neither by armchair philosophizing, nor by non-experimental linguistics, we can construct an idealized, normative model of conversation, some of the inputs and outputs of which correspond to those of the real psychological process. In the idealized model, we start with a sentence uttered, paired with the semantic content assigned to it by our semantic theory. In the model, idealized speaker-hearers extract information from such a pair by explicitly representing conversational maxims, plus propositions representing common conversational knowledge. Conclusions about what has been asserted, implied, and suggested are rationally inferred using all of this information. We validate the model by showing that these conclusions match those of real speakers in real speech situations.

The result is a rational reconstruction, rather than a causal explanation, of our ability to use sentences in communicative situations. When we show that some part of the information conveyed by an utterance results from the application of conversational maxims – as opposed to the meaning of the sentence uttered – we are showing that such information is rationally extractable from the utterance, together with a defensibly austere conception of meaning, even if real, non-idealized speakers don't follow this route either consciously or unconsciously. On this conception, linguistic meaning is a kind of *least common denominator*. It is information conventionally associated with expressions that must be mastered by a rational agent – over and above the agent's ability to reason intelligently and efficiently – in order to communicate with other members of the linguistic community. The point is not heuristic, but constitutive. This is what meaning is. Consequently, no matter what idiosyncratic processes we speakers actually go through in interpreting utterances, the question of what our sentences mean -- and of what part of that which is asserted and conveyed is due to meaning, and what part is due to other, pragmatic factors - is something to be determined by idealized rational reconstruction, not psycholinguistic research.

The point is illustrated by an elementary thought experiment. Suppose we could teach an ideally rational agent to interpret utterances by associating austere semantic contents with sentences, and explicitly employing Gricean and other rationally justifiable pragmatic rules. Suppose further that the judgments of this agent about what was asserted, implicated, and the like, matched those of competent speakers of L. If this were so, the agent would count as understanding L, and the success of our reconstruction would be evidence that the semantic and pragmatic theories it contained were correct. This would be so, whether or not the online psychological processing of ordinary speakers matched that given by our rational reconstruction. All we would know about the psychologically real processes of ordinary speakers would be that, for the most part, they tracked the rational idealization, even if they didn't match it point by point.

I suspect that one way in which the idealized model differs from whatever is really at work psychologically involves the model's use of semantic content as input to the process of interpreting utterances. I am not sure that this is a feature of our own cognitive architecture, or that our linguistic competence requires us to have psychologically robust representations which carry all and only the information semantically encoded by our sentences (while excluding all other information). What is required, I think, is simply that there be substantial uniformity among speakers and hearers regarding the salient information extracted from utterances in various contexts. From this perspective, the meaning of a sentence is an abstraction from the information asserted and conveyed by literal uses of it in normal contexts. It is that which is both (i) common to what is asserted and conveyed by utterances of it in all such contexts, and (ii) what an ideally rational agent would have to master – over and above the agent's ability to reason intelligently, and engage in cooperative social behavior – in order to communicate using the sentence with other speakers.^{xxiii} Understood in this way, sentence meaning is a notion with its own theoretical interest. However, it is not obvious that it is something individual speakers need to separate out in order to reliably identify what is asserted and conveyed.^{xxiv}

ⁱ One complication not covered by this definition involves cases in which one implicates q, not by saying (asserting) p, but by "making as if to say p" – i.e. by uttering something which, were it not for the presumption that the maxims were being obeyed, would count as saying p. Examples (discussed below) involve irony and metaphor. ⁱⁱ Paul Grice (1989), *Studies in the Way of Words*, (Cambridge: Harvard).

ⁱⁱⁱ Like Grice, I use the indirect-discourse sense of 'say' in such a way that *A says that S* is essentially equivalent to *A asserts that S*. For both Grice and me, saying/asserting that S is a way of committing oneself to the truth of the claim that S, distinct from merely implicating that S. In discussing implicature he typically contrasts "what is said by 'S'" with what an utterance of 'S' merely implicates – making it sound as if the sentence itself says or asserts something. What he means, I take it, is (i) that what a speaker says/asserts by uttering 'S' is to be contrasted with what the utterance merely implicates, and (ii) that what the speaker says/asserts is (in standard cases) to be identified with what is now called the semantic content of 'S' in the context – a statement determined by disambiguating 'S',

and settling the reference of any indexicals, time references, or other context-sensitive elements. ('S' is here used as a schematic letter, rather than a metalinguistic variable.)

^{iv} For an interesting discussion, see Richard Moran (1996), "Metaphor," in Hale and Wright, eds., *A Companion to the Philosophy of Language* (Oxford: Blackwell).

^v There is, I think, a tension in Grice's own discussions on this point. On the one hand, the official story he tells of what conversation implicatures are makes no room for implicatures that are said or asserted. On the other hand, some of his examples – e.g. of utterances of 'John took off his pants and he went to bed' as conversationally implicating that he took of his pants *and then* he went to bed – seem, quite clearly, to involve cases in which what he claims to be implicated *is* asserted by the speaker. This tension is an instance of the larger problem I will be addressing.

^{vi} Gilbert Ryle voices general worries of this sort – not specifically about disjunction but about all logical words – in chapter 8 of Ryle (1954), *Dilemmas* (Cambridge: Cambridge University Press).

^{vii} Thanks to Brian Bowman for raising this point.

^{viii} Note the difference between (i) *I have a child*, and (ii) *I have a son* as answers to *How many children do you have?*. While an utterance of (i) seems uncooperative, and a mildly puzzling, an utterance of (ii), in providing some extra information, sounds more natural, and more clearly suggests that I have only one child (even though the suggestion is still not part of what is asserted). Thanks for this to Nate Gadd.

^{ix} The contrast noted here between uses of instances of (3) to say (assert) that one has at least n F's, and uses of those instances to say (assert) that one has exactly n F's is noted on pp. 46-7 of Robyn Carston (1988), "Implicature, Explicature, and Truth-Theoretic Semantics," in Ruth Kempson, ed., *Mental Representation: The Interface Between Language and Reality*, (Cambridge: Cambridge University Press), and also on pp. 170-5 of Larry Horn (1992), "The Said and the Unsaid," in Chris Barker and David Dowty, eds., *Salt II: Proceedings of the Second Conference on Semantics and Linguistic Theory*, (Columbus Ohio: The Ohio State University).

^x These uses of numerals must be distinguished from those in which numerals function as simple singular terms standing for numbers – as in 6 + 3 = 9". The proposals under investigation in this paper only involve adjectival occurrences of numerals in quantified phrases.

xi Thanks to John MacFarlane, and to an anonymous referee, for pointing out the need to consider this reading.

^{xii} I leave it open how the * readings are generated, and what logical forms the relevant sentences have. My claim is only that these sentences can be used to assert propositions with the indicated truth conditions. On the ambiguity theory, this must be treated as a semantic fact.

^{xiii} Carston (1988) and Horn (1992) cite (i) – (iii) as cases in which assertive utterances result in "at most n" assertions.

(i) She can have 2000 calories a day without putting on weight.

(ii) The council houses are big enough for families with three kids.

(iii) Arnie is capable of shooting 70 on this golf course.

In my opinion, (iii) is out of place. There, '70' functions as a singular term rather than an adjective in a quantifier phrase. Needing no modification, it simply designates a score Arnie can shoot. (i) and (ii) may be useable to make "at most n" assertions, but they can also make "up to n" assertions. The two aren't distinguished in Carston and Horn.

^{xiv} In special cases the setting of k may require fine-tuning. Consider a dialog involving A, B, and C: (A) *Matriculated students are required to take at least three courses.* (B) *And they are allowed to take up to five.* (C) *Actually, A, they are required to take at least four courses. Perhaps you have forgotten the PE requirement.* Suppose C is right. Did B say something false? Though there are, I think, considerations on both sides, the principle of charity exerts considerable pressure to interpret (B) in a way that allows us to say "no". Resetting k to equal 4 will accommodate this pressure. (Another possibility would be to reinterpret B by assigning the narrow-scope interpretation to B's sentence, though the cost of doing so might be to weaken B's statement too much.) The principles governing the contextual interpretation, and reinterpretation, of utterances are subtle, highly complex, and in need of detailed examination. Although there are challenging issues to be resolved, they do not, in my view, undermine the plausibility of the overall picture. (Thanks to a referee for bringing this dialog to my attention.) ^{xv} This doesn't mean that no sentences containing numerals express complete propositions. When numerals occur as simple singular terms, as in 6+3 = 9 there is no incompleteness – nor is there incompleteness in 6 Fs + 3 Fs = 9 Fs. ^{xvi} Carston and Horn argue that sentences containing numerals should be given the kind of bare semantic analysis indicated here, with conversational facts and Gricean maxims contributing to enriched assertions. Carston says, "This is the line I would wish to pursue, taking all the numerals as having a single sense, neither an 'at least', an 'at most' nor an 'exactly' sense, these being determined pragmatically at the level of explicit content." (1988, p. 46). However, whereas she suggests extending the analysis to cases involving scalar implicatures generally, Horn (1992) argues persuasively against doing so. See his discussion for relevant data, plus an informative synopsis of the insights of earlier pragmatic literature -- including Jay Atlas's pithy observation that "only in the context of an NP does a numeral modifier have a meaning," in Atlas (1979), "How linguistics matters to philosophy," in Oh and Dineen, eds., *Syntax and Semantics 11: Presupposition*, New York: Academic Press, 1979.

^{xvii} Cases of obligatory enrichment are discussed in Soames (2005a), "Naming and Asserting," in Z. Szabo, ed., *Semantics vs. Pragmatics* (Oxford: Clarendon Press), where it is noted that the semantic contents of sentences containing possessive noun phrases only weakly constrain, without determining, the asserted relationship between the denotation of the possessor NP and the denotation of the possessed NP. See also Soames (forthcoming) "The Gap Between Meaning and Assertion," in Hackl, M. and R. Thornton, eds., *Asserting, Meaning, and Implying* (Oxford: Oxford University Press), where it is argued that some descriptions require contextually supplied temporal information, without which the semantic contents of sentences containing them are non-truth-evaluable. Cases of optional enrichment involving names, definite descriptions, and propositional attitude ascriptions are discussed in Soames (2005a), and Soames (2005b), "Why Incomplete Descriptions do not Defeat Russell's Theory of Descriptions," *Teorema*, Vol. XXIV, No. 3, 7-30.

^{xviii} Sometimes optional enrichment – though not required to obtain a complete proposition – is needed to make the speaker's remark relevant and informative. Two examples given on p. 39 of Carston (1988), are: *The park is some distance from where I live*, and *It will take us some time to get there*. The distinction between cases in which enrichment is required to secure a complete proposition, and those in which it is optional, parallels Kent Bach's distinction between *conversational implicitures* that complete semantically underdetermined sentences and those that merely expand propositions that are semantically expressed. See his path-breaking articles, Bach (1994), "Conversational Impliciture," *Mind and Language* 9, 124-62; and Bach (2001), "You don't say?," *Synthese*, 127, 11-31. An important early discussion of semantic underdetermination is Jay Atlas (1989), *Philosophy Without Ambiguity*, (Oxford: Clarendon Press).

^{xix} Similar explanations can be given for the examples in (7). In general, there is a wide range of possible enrichments of bare numerical quantifiers – which is not to say that for every conceivable enrichment there are

conversations which call for it. Far from it, discovering natural constraints on enrichment is a central task for future research.

^{xx} Had I answered, *San Francisco is 396 miles from here*, the situation would have been different. Having avoided a round number, I would naturally be understood as making a (false) claim about the exact distance, rather than a true claim of approximate distance. Horn (1992) makes a related point on page 173, which also contains references to other discussions of what he aptly terms "the pragmatics of approximation."

^{xxi} Although there is substantial evidence in favor of this natural account, there are also anomalies that remain unexplained. An example, brought to my attention by Timothy Williamson, is (i).

(ia) How many miles is it to San Francisco?

(ib) *Its 400 miles to San Francisco, but not more that 398.

(ic) Its approximately 400 miles to San Francisco, but not more than 398.

Although (ic) could be uttered in reply to (ia), (ib) could not. If enrichment of '400 miles' is required, and 'approximately 400 miles' is generally a candidate, why isn't it available here? Its not clear. Perhaps, this approach to "loose talk" is wrong, and 'approximately' isn't generally available. However, (ii) suggests that something else may be at work.

(iia) How many beers do you have in the fridge?

(iib) ?* I have 3 beers in the fridge, but not more than 6.

(iic) I have at least 3 beers in the fridge, but not more than 6.

Although (iic) could be uttered in reply to (iia), (iib) is highly questionable. Yet we know that '3 beers' can often be understood as 'at least 3 beers' (either due to pragmatic enrichment or to semantic ambiguity). 'How many' questions seem to invite 'exactly' understandings. What else may be going on in cases like these deserves further exploration. ^{xxii} See Bach (1994), 140-41 for similar conclusions. The distinction I draw between the role of Gricean maxims in contributing to assertions and their role in determining conversational implicatures also has a counterpart in Relevance Theory, a speculative, psychological theory about the causal processes involved in sentence comprehension descending from D. Sperber and D. Wilson (1986), *Relevance* (Oxford: Blackwell). In it, radically psychologized versions of Grice's maxims are rolled into a principle of cognitive dynamics that maximizes relevant information derived from an utterance while minimizing processing costs. Semantic information is taken to be the output of an automatic decoder in that process. Although this information is sometimes a complete proposition, often it isn't. Either way, it is seen as initiating an unconscious inference-drawing process that terminates with the propositions asserted by the utterance, when relevance has been maximized at minimal cost. When, as it often is, pragmatic information is incorporated in this outcome, the resulting (asserted) information is called *an explicature*. By contrast, an *implicature* is any other proposition communicated by an utterance, the content of which includes pragmatically inferred (non-semantic) material. For recent explications see Carston (2004), "Explicature and Semantics," in S. Davis and B. Gillon, eds., *Semantics*, (New York: Oxford), and (2002) "Linguistic Meaning, Communicated Meaning, and Cognitive Pragmatics," *Mind and Language* 17. From this point of view, the examples discussed here involving numerical quantifiers are called *explicatures*. Although there is substantial overlap between what I say and what relevance theorists say about assertions and implicatures in cases like these, our theoretical perspectives stand in stark contrast, as will become clear in the next section.

A similar comment applies to Recanati (2003), *Literal Meaning*, (Cambridge: Cambridge University Press). Like Bach, the relevance theorists, and me -- he takes what a speaker means and asserts to often contain pragmatic information, over and above the semantic content of the sentence uttered. Also like us, he sees conversational implicatures as pragmatic information that is communicated, beyond what the speaker asserts. His model of implicature is, with one caveat, essentially what one would expect from Grice himself -- if one substituted his pragmatic notion of what is asserted by the speaker for Grice's more overtly semantic notion of *what is said (by the sentence)*, as that which generates implicatures. The caveat is that, like relevance theorists, but unlike Grice and me, Recanati situates pragmatics within a psychological theory of the causal mechanisms involved in speech comprehension. Other distinctive features of his position are (i) a sharp distinction between (a) unconscious pragmatic processes that generate what an utterance asserts and (b) potentially conscious processes that generate conversational implicatures, and involve tacit appeal to conversational maxims, (ii) a seeming rejection of the notion of the semantic content of a sentence (as opposed to the lexical meanings of its words) as playing any role in pragmatics, and (iii) a remarkably unconstrained conception of the pragmatic enrichment, transformation, or deletion of the meager semantic information involved in determining what an utterance asserts.

(2005a) and (forthcoming). As for (ii), when a piece of information M is the meaning of S, both ideal agents and ordinary speakers will, in some way, associate M with S. Nevertheless, ordinary speakers may fail to recognize M

as the meaning of S, because their competence with S doesn't require any psychologically significant distinction between M and other, nonsemantic, information they associate with S.

^{xxiv} Thanks to Kent Bach, Chris Barker, Brian Bowman, Nate Gadd, Jeff King, John MacFarlane, and two referees for their helpful comments. Special thanks to Kent for sending an early draft to Larry Horn, who informed me of the early championing, in Carston (1988) and Horn (1992), of a view about cardinal numerals similar to the one defended here.