

What is said and the semantics/pragmatics distinction

François Recanati

► **To cite this version:**

François Recanati. What is said and the semantics/pragmatics distinction. WOC 2002: Semantics and Pragmatics, Oct 2002, Genoa, Italy, 2002. <ijn_00000272>

HAL Id: ijn_00000272

https://jeannicod.ccsd.cnrs.fr/ijn_00000272

Submitted on 4 Nov 2002

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.

'What is said' and the semantics/pragmatics distinction

François Recanati

Institut Jean Nicod

*Workshop on Context 2002:
Semantics vs. Pragmatics
Genoa, October 25 & 26 2002*

I.

It is customary in pragmatics to ascribe to an utterance literal truth-conditions at variance with the intuitive truth-conditions which the conversational participants themselves would ascribe to that utterance. For example, the proposition literally expressed by

(1) I have three children

is standardly taken to be the proposition that the speaker has at least three children, i.e., no less than three but possibly more. In certain contexts this corresponds to what the speaker actually means (as when I say, 'If I have three children I can benefit from lower rates on public transportation') but in other contexts what the speaker means is quite different. Suppose for example that I am asked how many children I have and that I reply by uttering (1). Clearly, in this context, I mean that I have (exactly) three children — no more and no less. This is standardly accounted for by saying that the proposition literally expressed, to the effect that I have at least three children, combines with the 'implicature' that I have no more than three children (a generalized implicature that is accounted for in terms of the maxim of quantity);¹ as a result of

¹ As Grice puts in in one of his early papers, "one should not make a weaker statement rather than a stronger one unless there is a good reason for so doing" (Grice 1961: 132). Since the statement that one has (at least) three children is weaker than the statement that one has n children (for $n > 3$), the maxim is obeyed only if the speaker has no more than three children.

this combination, what is globally communicated — and what I actually mean — is the proposition that I have exactly three children. Now *this is the only proposition I am conscious of expressing by my utterance*; in particular, I am unaware of having expressed the 'minimal' proposition that I have at least three children. To account for this obvious fact, the theorist claims that we are aware only of what is globally conveyed or 'communicated' by the utterance. Analysing this into 'what is literally said' and 'what is implied' is the linguist's task, not something that is incumbent upon the normal language user. Figure 1 illustrates this widespread conception.

[Figure 1 about here]

One problem with this conception is that it lacks generality. It turns out that there are two sorts of case. On the one hand there are prototypical cases of implied meaning, in which the participants in the speech situation are aware both of what is said and of what is implied, and also of the inferential connection between them. On the other hand, there are the cases illustrated by (1). Given his willingness to treat certain aspects of the intuitive meaning of (1) as conversational implicatures external to what is literally said, the theorist must explain why those implicatures, unlike the prototypical cases, do not have the property of conscious 'availability'.

The only explanation I have come across in the literature makes use of Grice's distinction between 'generalized' and 'particularized' conversational implicatures, i.e. between implicatures which arise 'by default', without any particular context or special scenario being necessary, and those which require such specific contexts. In contrast to the latter, the former are 'hard to distinguish from the *semantic* content of linguistic expressions, because such implicatures [are] routinely associated with linguistic expressions in all ordinary contexts' (Levinson 1983: 127). Generalized implicatures are unconsciously and automatically generated and interpreted. They belong to the 'micropragmatic' rather than to the 'macropragmatic' level, in Robin Campbell's typology:

(If the speaker has more than three children, he should have made a stronger statement.) The statement 'I have three children' therefore implicates that the speaker has no more than 3

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)

But there are problems with this explanation. According to Horn (1992), the generalized nature of an implicature does not entail its conscious unavailability — its 'cryptic' character. In other words, it is possible for an implicature to be both 'generalized' and intuitively accessible as an implicature distinct from what is said. Thus Horn insists that the generalized scalar implicature from 'some' to 'not all' *is* consciously available (in contrast to that from 'three' to 'exactly three'). A speaker saying 'Some students came to the meeting' normally implies that not all students came, and when this is so there is (Horn claims) no tendency on the part of the interpreter to conflate the implicature with what is said. This is actually debatable, for the 'implicature' at issue can arise at sub-sentential level (e.g. 'He believes some students came'), and in such cases there are reasons to doubt that the availability condition is satisfied. Be that as it may, the 'generalization' of an implicature does not seem to be necessary for its unconscious character. Many particularized 'bridging' inferences are automatic and unconscious. To take an example from Robyn Carston, 'He went to the cliff and jumped' is readily interpreted as saying that the person referred to jumped *over the cliff*, even though this is only contextually suggested.

II.

In earlier writings (Recanati 1989, 1993) I put forward a conception diametrically opposed to that illustrated by Figure 1 above. 'What is said', I held, *is* consciously available to the participants in the speech situation (Figure 2).

[Figure 2 about here]

In this framework 'what is communicated' is not a distinct level where 'what is said'

children.

and 'what is implied' have been merged and integrated into a unified whole; it is merely a name for the level at which we find both what is said and what is implied, which level is characterized by conscious accessibility. On this picture, there are only two basic levels: the bottom level in which we find both the meaning of the sentence and the contextual factors which combine with it to yield what is said; and the top level at which we find both what is said and what is implied, both being consciously accessible (and accessible as distinct).

The availability of what is said follows from Grice's idea that saying itself is a variety of nonnatural meaning. One of the distinguishing characteristics of nonnatural meaning, on Grice's analysis, is its essential overtness. Nonnatural meaning works by openly letting the addressee recognize one's primary intention (e.g. the intention to impart a certain piece of information, or the intention to have the addressee behave in a certain way), that is, by (openly) expressing that intention so as to make it graspable. This can be done in all sorts of ways, verbal or nonverbal. Even if we restrict ourselves to verbal communication, there are many ways in which we can mean things by uttering words. *Saying* is one way ; *implying* is another one.

The view that 'saying' is a variety of nonnatural meaning entails that what is said (like what is meant in general, including what is implied) *must* be available — it must be open to public view. That is so because nonnatural meaning is essentially a matter of intention-recognition. On this view what is said by uttering a sentence depends upon, and can hardly be severed from, the speaker's publicly recognizable intentions. Hence my 'Availability Principle' (Recanati 1993: 248), according to which 'what is said' must be analysed in conformity to the intuitions shared by those who fully understand the utterance — typically the speaker and the hearer, in a normal conversational setting.

I take the conversational participants' intuitions concerning what is said to be revealed by their views concerning the utterance's truth-conditions. I assume that whoever fully understands a declarative utterance knows which state of affairs would possibly constitute a truth-maker for that utterance, i.e. knows in what sort of circumstance it would be true. The ability to pair an utterance with a type of situation in this way is more basic than, and in any case does not presuppose, the ability to *report* what is said by using indirect speech; it does not even presuppose mastery of the notion of 'saying'. Thus the proper way to elicit such intuitions is not to ask the subjects 'What do you think is said (as opposed to implied or whatever) by this

sentence as uttered in that situation'? I therefore agree with Bach's criticism of the experiments through which Gibbs and Moise (1997) attempted to support my Availability-based approach.² However, Bach himself uses what he calls the 'IQ test' to determine what is said, that is, he ties what is said to indirect speech reports of what is said (Bach 1994a: 278, 1999, 2001). I find this procedure most objectionable, and that is *not* what I mean when I claim that what is said should be individuated according to the intuitions of normal interpreters. Thus I strongly disagree with Cappelen and Lepore's surprising statement:

We ourselves don't see how to elicit intuitions about what-is-said by an utterance of a sentence without appealing to intuitions about the accuracy of indirect reports of the form 'He said that...' or 'What he said is that...' or even 'What was said is that...' (Cappelen and Lepore, 'On an alleged connection between indirect speech and the theory of meaning', *Mind and Language* 12: 278-296, 1997, p.280)

I find this statement surprising, because there obviously *is* another way of eliciting truth-conditional intuitions. One has simply to provide subjects with scenarios describing situations, or, even better, with — possibly animated — pictures of situations, and to ask them to evaluate the target utterance as true or false with respect to the situations in question. That procedure has been used by several researchers to test speaker's intuitions about e.g. the truth-conditions of donkey sentences. Thus Bart Geurts describes his experimental set-up (inspired from earlier work by Yoon) as follows:

Twenty native speakers of Dutch were asked to judge whether or not donkey sentences correctly described pictured situations. Instructions urged subjects to answer either true or false, but they were also given the option of leaving the matter open in case they couldn't make up their minds. (B. Geurts, 'Donkey Business', *Linguistics and Philosophy* 25: 129-156, 2002, p. 135)

This procedure presupposes that normal interpreters have intuitions concerning the

² K. Bach, 'You don't say', *Synthese* 128: 15-44, 2001, p. 43.

truth-conditional content of utterances. On my view, those intuitions correspond to a certain 'level' in the comprehension process — a level that a proper theory of language understanding must capture. That is the level of 'what is said'.

III.

From a psychological point of view, we can draw a helpful parallel between understanding what one is told and understanding what one sees. In vision, the retinal stimuli undergo a complex (multistage) train of processing which ultimately outputs a conscious perception, with the dual character noted by Brentano : the subject is aware both of what he sees, and of the fact that he is seeing it. Although more complex in certain respects, the situation with language is similar. The auditory signal undergoes a multistage train of processing which ultimately outputs a conceptual experience : the subject understands what is said. This is very much like (high-level) perception. If I am told that it is four o'clock, I hear that it is four o'clock, just as, when I look at my watch, I see that it is four o'clock. Like the visual experience, the locutionary experience possesses a dual character : we are aware both of what is said, and of the fact that the speaker is saying it.

In calling understanding an *experience*, like perception, I want to stress its conscious character. Understanding what is said involves entertaining a mental representation of the subject-matter of the utterance that is both determinate enough (truth-evaluable) *and* consciously available to the subject. This suggests a new criterion for demarcating what is said. Instead of looking at things from the linguistic side and equating 'what is said' with the minimal proposition one arrives at through 'saturation', as semanticists usually do, we can take a more psychological stance and equate what is said with (the semantic content of) the conscious output of the complex train of processing which underlies comprehension.³

To be sure, that output itself is subject to further processing through e.g. inferential exploitation. Consider, once again, vision. Seeing John's car, I can infer that he is around. Similarly, hearing John say that it is late, I can infer that he wants

me to leave. Just as what is seen corresponds to the primary conscious output of visual processing, not to what can be secondarily derived from it, 'what is said' corresponds to the primary truth-evaluable representation made available to the subject (at the personal level) as a result of processing the sentence.

Accordingly, I distinguish between two sorts of pragmatic process. The contextual processes which are (subpersonally) involved in the determination of what is said I call *primary* pragmatic processes. In contrast, *secondary* pragmatic processes are ordinary inferential processes taking us from what is said, or rather from the speaker's saying of what is said, to something that (under standard assumptions of rationality and cooperativeness) follows from the fact that the speaker has said what she has said. To the extent that the speaker overtly intends the hearer to recognize such consequences as following from her speech act, they form an integral part of what the speaker means by her utterance. That is, roughly, Grice's theory of 'conversational implicature'. An essential aspect of that theory is that the hearer must be able to recognize what is said and to work out the inferential connection between what is said and what is implied by saying it. Again, it follows that what is said must be consciously available to the interpreter. It must satisfy what I call the Availability constraint.

IV.

The notion of 'what is said' we end up with by following this route is different from the standard notion used in semantics: it is not as close to the linguistic meaning of the sentence. To get what is said in the standard, semantic sense, we must assign semantic values to indexicals and free variables in logical form. That process, which I call 'saturation', is the only pragmatic process that is allowed to affect what is said in the standard semantic sense. Beside saturation, however, there are other pragmatic processes which contribute to shaping what is said in the psychological sense. In contrast to saturation, which is linguistically mandated (bottom up), the other pragmatic processes are optional and context-driven (top-down), so they are

3 As Ian Rumfitt once put it, "what is said in the course of an utterance is nothing other than what somebody who understands the utterance understands to be said" ("Content and Context:

construed as semantically irrelevant. The paradigm cases is free enrichment, illustrated by example (2):

(2) Mary took out her key and opened the door

In virtue of a 'bridging inference', we naturally understand the second conjunct as meaning that Mary opened the door with the key mentioned in the first conjunct ; yet this is not explicitly articulated in the sentence.

In typical cases free enrichment consists in making the interpretation of some expression in the sentence contextually more specific. This process has sometimes been described in the literature as 'specification'. For example the mass term 'rabbit' will be preferentially interpreted as meaning *rabbit fur* in the context of 'He wears rabbit' and as meaning *rabbit meat* in the context of 'He eats rabbit' (Nunberg and Zaenen 1992).⁴ This not a matter of selecting a particular value in a finite set; with a little imagination, one can think of dozens of possible interpretations for 'rabbit' by manipulating the stipulated context of utterance; and there is no limit to the number of interpretations one can imagine in such a way. Nor can the process of specification be construed as linguistically mandated, that is, as involving a hidden variable. Were it linguistically mandated (bottom up), it would be mandatory, but it is not: In some contexts the mass term 'rabbit' means nothing more than RABBIT STUFF ('after the accident, there was rabbit all over the highway').

The converse of enrichment is loosening (Cohen 19xx, Carston 1996). There is loosening whenever a condition of application packed into the concept literally expressed by a predicate is contextually dropped so that the application of the predicate is widened. An example is 'The ATM swallowed my credit card'. There can be no real swallowing on the part of an ATM, since ATMs are not living organisms with the right bodily equipment for swallowing. By relaxing the conditions of application for 'swallow', we construct an ad hoc concept with wider application

A third type of primary pragmatic process that is not linguistically mandated (bottom up) but contextually driven is semantic transfer (Nunberg 1979, 1995). In

the Paratactic Theory Revisited and Revised", *Mind* 102: 429-453, at p. 439).

⁴ The term 'specification' comes from Bach. Nunberg & Zaenen use 'precision' in the same sense.

transfer the output is neither an enriched nor an impoverished version of the concept literally expressed by the input expression. It's a different concept altogether, bearing a systematic relation to it. Thus 'parked out back' denotes either the property a car has when it is parked out back, or a different property, namely the property a car-owner has whenever his or her car has the former property ('I am parked out back'). Arguably, 'parked out back' literally denotes the former property, and comes to denote the latter property as a result of transfer (Nunberg 1995). Similarly, the expression 'ham sandwich' in 'The ham sandwich left without paying' arguably denotes, through transfer, the derived property HAM_SANDWICH_ORDERER rather than the linguistically encoded property HAM_SANDWICH.

Despite their optional character, pragmatic processes such as enrichment, loosening and transfer affect the intuitive truth-conditions of utterances. If we use the availability criterion to demarcate what is said, as I suggest, then such processes must be treated as primary rather than secondary. And the same consideration applies to the so-called generalized conversational implicature which is responsible for the 'exactly' reading of the numeral in example (1). In this way, we solve the difficulty raised in section 1. We no longer have two sorts of case of implicature — the prototypical cases where the interlocutors are aware of what is said, aware of what is implied, and aware of the inferential connection between them, and the cases in which there is no such awareness. Conscious awareness is now a built-in feature of both what is said and the implicatures. That is so because what is said is the conscious output of linguistic-cum-pragmatic processing, and the implicatures correspond to further conscious representations inferentially derived, at the personal rather than sub-personal level, from what is said (or, rather, from the speaker's saying what is said). The alleged cases in which the speech participants themselves are not distinctly aware of what is said and of what is implied are reclassified: they are no longer treated as cases of 'implicature', strictly speaking, but as cases in which a primary pragmatic process operates in the (sub-personal) determination of what is said.

V.

So far I have followed Grice, who construes saying as a variety of meaning. But this pragmatic approach to 'saying' is controversial. Most philosophers use the notion of

'what is said' (or 'the proposition expressed') in such a way that it is *not* a 'pragmatic' notion — having to do with what the speaker means or with what the hearer understands. What is said is supposed to be a property of the *sentence* (with respect to the context at hand) — a property which it has in virtue of the rules of the language.

I will discuss the nonpragmatic construal of what is said in the next section. For the time being, I'm interested in the pragmatic construal, based on Grice's idea, and the alleged objections it raises.

The first objection is this. If, following Grice, we construe saying as a variety of meaning, we will be prevented from acknowledging an important class of cases in which the speaker does not mean what he says. Irony is a good example of that class of cases. If I say 'John is a fine friend' ironically, in a context in which it is obvious to everybody that I think just the opposite, it is clear that I do not mean what I say: I mean the opposite. Still, I say that John is a fine friend. Grice's construal of saying as a variety of meaning prevents him from acknowledging that fact. According to Grice, when I say 'John is a fine friend' in the mentioned situation, I do not *really* say that John is a fine friend — I *pretend* to be saying it. The pragmatic construal of saying forces Grice to draw a distinction between 'saying' and 'making as if to say'.

As far as I am concerned, I find Grice's distinction (between genuine saying and making as if to say) perfectly legitimate, but I can understand the worries of those who feel that the notion of 'saying' he uses is too much on the pragmatic side. We certainly need a notion of 'what is said' which captures *the objective content of an utterance irrespective of its pragmatic force as a serious assertion or as an ironical utterance*. Still, I find the objection superficial, for it is quite easy to actually construct the desired notion within Grice's own framework. Grice uses 'say' in a strict sense. In that sense whatever is said must be meant. But we can easily define a broader sense for 'say':

S says that p, in the broad sense, iff he either says that p (in the strict sense) or makes as if to say that p (again, in the strict sense of 'say').

I will henceforth use 'say' in that broad sense, which remains within the confines of the pragmatic construal.

Another objection to the pragmatic construal focusses on the loss of objectivity

that allegedly goes with it. What is said is objective in the sense that it is possible both for the speaker to make a mistake and say something other than what he means, and for the hearer to misunderstand what the speaker is saying. Those mistakes are possible, the objector will argue, because what is said is an objective property of the sentence (in context). But on the pragmatic construal, it is not clear that this objectivity can be captured. Imagine the following situation: the speaker wants to say that Paul is tall, and, mistaking Tim for Paul, says 'He is tall' while pointing to Tim. The speaker thus inadvertently says that Tim is tall. Now imagine that the hearer also mistakes Tim for Paul. Thanks to this lucky mistake, he grasps what the speaker means, thinking that this is what he has said. The speaker and the hearer therefore converge on a certain interpretation, which is not objectively what was said, but which they both (mistakenly) think is what was said. How, in the framework I have sketched, will it be possible to dissociate what is actually said from the protagonists' mistaken apprehension of what is said? Have we not equated what is said with their understanding of what is said?

We have not. We have equated what is said with what a *normal interpreter* would understand as being said, in the context at hand. A normal interpreter knows which sentence was uttered, knows the meaning of that sentence, knows the relevant contextual facts (who is being pointed to, etc.).⁵ Ordinary users of the language *are* normal interpreters, in most situations. They know the relevant facts and have the relevant abilities. But there are situations (as in the above example) where the actual users make mistakes and are not normal interpreters. In such situations their interpretations do not fix what is said. To determine what is said, we need to look at the interpretation that a normal interpreter would give. This is objective enough, yet remains within the confines of the pragmatic construal.

VI.

I have presented the availability-based approach as an alternative to the standard view. But why not combine them? Many authors hold that *there are two equally legitimate notions of what is said*: a purely semantic notion, corresponding to the

'minimal' proposition resulting from saturation, and a pragmatic notion ('what is stated' as opposed to what is implied). Only what is said in the pragmatic sense needs to satisfy the availability constraint. So one can maintain that (1) expresses the proposition that the speaker has at least three children, while conceding that what the speaker states is that he has exactly three children. On this picture saturation maps the linguistic meaning of a sentence to the minimal proposition literally expressed, while primary pragmatic processes of the optional variety map that proposition to that which is actually asserted (Figure 3).

[Figure 3 about here]

This view – which I call the Syncretic View – has been argued for by many authors, including several members of this audience. A recent statement can be found in Scott Soames' book *Beyond Rigidity*. He writes:

When speaking of the information carried by an assertive utterance of a sentence in a context, one must distinguish (i) the semantic content of the sentence uttered in the context; (ii) what the speaker says (asserts) by uttering the sentence; (iii) what the speaker implies, implicates, or suggests... (i) is standardly included in (ii), but... in the case of many utterances, (ii) is not exhausted by (i). (Soames 2002: 86)

In Recanati 2001 I have criticized the Syncretic View, on the following grounds. What is said, in the purely semantic sense, is supposed to be determined by the rules of the language (with respect to the context) quite independent of speaker's meaning : it is what *the sentence* says as opposed to what *the speaker* says. As Bach points out, what is said, in the semantic sense, "excludes anything that is determined by [the speaker's] communicative intention (if it included that, then what is said would be partly a pragmatic matter)" (Bach 2001: 21). Precisely for that reason, I claim that there is no such thing : no such thing as a complete proposition autonomously determined by the rules of the language with respect to the context but

5 This is all tacit knowledge, not the sort of 'conscious awareness' I talk about in connection

independent of speaker's meaning. In order to reach a complete proposition through saturation, we *must* appeal to speaker's meaning. That is the lesson of semantic underdeterminacy. Thus the demonstrative pronoun 'he' refers to the male person whom the speaker who utters the demonstrative refers to in uttering it. No semantic reference without speaker's reference, in such a case. Because of the well-documented phenomenon of semantic underdeterminacy, there is no such thing as 'what the sentence says'.

There is another possible interpretation for the semantic notion of 'what is said', however. Instead of construing what is said as a nonpragmatic property of the sentence, independent of speaker's meaning, we can start with the pragmatic notion of what is said, and define the semantic notion in terms of it. What is said in the minimal sense can thus be defined as what is said in the full-fledged, pragmatic sense *minus* the contextual ingredients that are optional and whose provision is context-driven. To filter out the optional ingredients, while retaining the contextual ingredients that are necessary for propositionality (reference of indexicals, etc.), one may follow Soames and define the semantic content of a sentence *s* relative to a context *C* as

that which would be asserted and conveyed by an assertive utterance of *s* in any normal context in which the reference of all indexicals in *s* is the same as their reference in *C*. (Soames 2002: 106).

Soames's strategy therefore consists in *abstracting* what is said in the semantic sense *from* what is said in the pragmatic sense. That is possible because, according to Soames, the semantic content of the sentence is *included in* the content of the assertion. To get to the semantic content, one only has to filter out those aspects of assertion content that go beyond semantic content and are tied to specific contexts of utterance. What remains, i.e. the 'common denominator', is the minimal proposition expressed by the sentence itself:

The semantic content of a sentence relative to a context is information that a competent speaker/hearer can confidently take to be asserted and conveyed

with secondary pragmatic processes.

by an utterance of the sentence in the context, no matter what else may be asserted, conveyed, or imparted. It is a sort of minimal common denominator determined by the linguistic knowledge shared by all competent speakers, together with contextually relevant facts such as the time, place, and agent of the context; the identity of individuals demonstrated by the speaker; and the referents of the names, as used in the context. As such, the semantic content of a sentence functions as a sort of minimal core around which speaker/hearers can structure the totality of information the sentence is used to communicate in a given context. (Soames 2002: 109)

This alternative strategy is also a failure, however. What is said in the minimal, semantic sense cannot be abstracted from what is said in the pragmatic sense simply because *it need not be part of it*. Soames's claim that semantic content is included in assertion content seems plausible because his examples are all cases in which the asserted content is richer than the alleged semantic content. He gives the following sort of example:

A man goes into a coffee shop and sits at the counter. The waitress asks him what he wants. He says, "I would like coffee, please." The sentence uttered is unspecific in several respects — its semantic content does not indicate whether the coffee is to be in form of beans, grounds, or liquid, nor does it indicate whether the amount in question is a drop, a cup, a gallon, a sack, or a barrel. Nevertheless, it is obvious from the situation what the man has in mind, and the waitress is in no doubt about what to do. She brings him a cup of freshly brewed coffee. If asked to describe the transaction, she might well say, "He ordered a cup of coffee" or "He said he wanted a cup of coffee", meaning, of course, the brewed, drinkable kind. In so doing, she would, quite correctly, be reporting the content of the man's order, or assertion, as going beyond the semantic content of the sentence he uttered. (Soames 2002: 78)

Now Soames thinks that in such cases several propositions are asserted, including (i) the unspecific proposition literally expressed (to the effect that the man wants coffee in some form or other) and (ii) more specific propositions recoverable from the literal proposition and the context. Those more specific propositions resulting from

enrichment are tied to the particular context in which they are generated, hence they can be filtered out by considering other contexts in which that sentence would be uttered and the indexicals would be given the same semantic values. Soames equates the minimal proposition expressed by the sentence with the proposition which would be asserted in *all* such contexts.

The problem is that enrichment is only one pragmatic process among others. Beside enrichment, there are other primary pragmatic processes, like loosening or semantic transfer, which are optional and take us away from the 'minimal' proposition allegedly expressed by the sentence. Only in the case of enrichment, however, is it plausible to suggest that the minimal proposition itself is part of what is asserted.

Soames glosses assertion in terms of commitment: "assertively uttering a sentence with the intention to assert or convey p involves doing so with the intention of committing oneself to p " (73). Since one cannot commit oneself to the truth of a specific proposition p without committing oneself to the truth of a less specific proposition q which it entails, it makes sense to say that the minimal proposition q is asserted when one (intuitively) asserts an enriched proposition p . But the principle that the minimal proposition is part of what is asserted (hence can be abstracted from it) does not hold when the primary processes at issue are instances of loosening or transfer. The speaker who assertively utters 'The ham sandwich left without paying', thereby referring to the ham sandwich orderer, does not assert the minimal proposition that the sandwich itself left without paying! Hence the minimal proposition cannot be defined as the common denominator — what is asserted in all contexts in which the sentence is uttered and the indexicals are given the same semantic values as in the current context.

To be sure, the counterexamples involving loosening or semantic transfer are taken care of, in Soame's framework, by the notion of a 'normal' context, that is, of a context in which the sentence "is used with its literal meaning", i.e. "is not used metaphorically, ironically, or sarcastically". However this qualification cannot be invoked in the context of the present debate without begging the question. We are supposed to start with the intuitive (pragmatic) notion of what is said, which sometimes *is* affected by loosening or transfer. If, following Soames, we want to build the notion of the minimal proposition out of what is said in that sense, we cannot arbitrarily set aside the cases that potentially threaten the enterprise.

VII.

The minimal proposition which the Syncretic View posits as the semantic content of the utterance, and which results from saturation alone, is not autonomously determined by the rules of the language independent of speaker's meaning. At the same time, the minimal proposition does not necessarily correspond to an aspect of what the speaker asserts. The minimal proposition is a hybrid which goes beyond what is determined by the rules of the language yet has no psychological reality and need not be entertained or represented at any point in the process of understanding the utterance (Recanati 1995, 2001).

Do we need such a notion in theorizing about language and communication? Many people in this audience and elsewhere claim that we do, but I can hardly understand why. In a recent paper ('Semantics, Pragmatics, and the Role of Semantic Content'), King and Stanley argue against semantic theories which ascribe (functional) characters to sentences on the basis of the characters of their parts, on the grounds that "the job of character is to give us content, and we can assign contents to complex expressions in contexts using only the characters of the parts, and combining the contents they determine in those contexts." They conclude that

Both a semantics that assigns characters to simple expressions and recursively assigns characters to complex expressions *and* a semantics that assigns characters to only simple expressions allow for an assignment of the same contents in contexts to simple and complex expressions. So unless the functional characters of complex expressions have some *additional* job to do, they are unnecessary. But there seems to be no such additional job.

The same sort of argument seems to me to rule out the minimal proposition as unnecessary. What must ultimately be accounted for is what speakers say in the pragmatic sense. The job of characters, contents etc. is to contribute to the overall explanation. But it is sufficient to assign semantic contents (in context) to simple expressions. Pragmatic processes will operate on those contents, and the composition rules will compose the resulting pragmatic values, thereby yielding the content of the speaker's assertion. Of course it is possible to let the composition rules compose the semantic contents of the constituent expressions, thereby yielding the

minimal proposition expressed by the sentence. (An absurd proposition, in certain cases.) However, the content of the speaker's assertion will still be determined by composing the pragmatic values resulting from the operation of pragmatic processes on the contents of the constituent expressions; so it is unclear what additional job the minimal proposition is supposed to be doing.

If one insists in using a purely semantic notion of 'what is said', i.e. a notion of what is said which is propositional (truth-evaluable) yet "excludes anything that is determined by the speaker's communicative intention", there is a much better candidate than the alleged minimal proposition. For every utterance, there arguably is a proposition which it expresses in virtue solely of the rules of the language, independent of speaker's meaning: that is the 'reflexive' proposition in the sense of John Perry (a variant of Stalnaker's diagonal proposition).⁶ The main difference between the minimal proposition and the reflexive proposition is that the reflexive proposition is determined *before* the process of saturation takes place. The reflexive proposition can't be determined unless the sentence is tokened, but no substantial knowledge of the context of utterance is required to determine it. Thus an utterance *u* of the sentence 'I am French' expresses the reflexive proposition that *the utterer of u is French*. That it does not presuppose saturation is precisely what makes the reflexive proposition useful, since in most cases saturation proceeds by appeal to speaker's meaning. If one wants a proposition that's determined on purely semantic grounds, one had better not have it depend upon the process of saturation.

Soames considers the possibility of equating 'what the sentence says' with the reflexive proposition or something close to it, but he rejects that option with the following argument:

Consider... the first-person singular pronoun as it occurs in a sentence *I am F*. There is no such thing as "what this sentence says" independent of the context of utterance in which it is used. The competence conditions associated with the first-person singular pronoun guarantee that when I assertively utter the sentence, I use it to say something about me, whereas when you assertively utter it, you use it to say something about you. One might be tempted to

suppose that there is some more general thing that the sentence "says" in every context — namely, the proposition expressed by *the speaker is F* (or some such thing). But this will not do. Our notion of "what a sentence says" is tied to what speakers who assertively utter the sentence say. Typically, when I assertively utter *I am F*, I don't assert that I am speaking or using language at all. Further, the proposition that I assert when I assertively utter such a sentence may be true in a possible circumstance in which no one is using language, and someone may believe this proposition without believing anything about me being a speaker. (Soames 2002: 104)

Soames' main objection is that the alleged reflexive proposition is not (part of) what the speaker asserts. As we have seen, however, the same thing often holds for the minimal proposition posited by the syncretists. The advantage of the reflexive proposition over the minimal proposition is that it (the reflexive proposition) is determined solely by the rules of the language, independent of speaker's meaning, in such a way that there *is* a path to the reflexive proposition that does not go through the speaker's meaning; hence it does not matter much if that proposition can't be reached by abstraction from what the speaker asserts.

The reflexive proposition is admittedly distinct from that which the speaker asserts — they have different possible worlds truth-conditions, as Soames points out — but why is this an objection? Are we not supposed to draw a distinction between *the proposition expressed by the sentence* and *the proposition asserted by the speaker who utters that sentence*? Note that we can, if we wish, incorporate into the reflexive proposition something to the effect that the linguistic mode of presentation associated with the first person pronoun will not be part of the proposition asserted, while the reference it contextually determines will be. Thus we might take the reflexive proposition expressed by an utterance *u* of 'I am French' to be the proposition that *there is an x such that x utters u and u is true iff x is French*. This comes as close as one can get to capturing, in propositional format, the information provided by the utterance in virtue solely of the linguistic meaning of the sentence 'I am French'. (See Recanati 1993 for an analysis along those lines.) Such a reflexive

6 See R. Stalnaker, *Context and Content*, Oxford University Press 1999; J. Perry, *Reference*

proposition determines that the proposition contextually asserted by 'I am French' will consist of the reference of 'I' and the property of being French. The reflexive proposition is therefore 'tied to what speakers who assertively utter the sentence say', even if it is not (part of) what they say.

I conclude that there may be a way of preserving the notion of 'what the sentence says', in the purely semantic sense, if one wants to; but it does not support the Syncretic View with its four levels (sentence meaning, what is said in the semantic sense, what is said in the pragmatic sense, and what is implied). What characterizes the reflexive proposition is that, although fully propositional, it does not incorporate those contextual ingredients whose provision is linguistically mandated; it is much closer to the linguistic meaning of the sentence — indeed it is directly and immediately determined by the linguistic meaning of the sentence. Appealing to the reflexive proposition instead of the minimal proposition takes us back to the availability-based approach with its three basic levels: the linguistic meaning of the sentence (and the reflexive proposition it directly and immediately determines); what is said in the pragmatic sense; and what is implied or otherwise conveyed by the utterance.