

Pragmatic Composition?

Peter Pagin

June 15, 2005

Abstract

Over the years, compositional semantics has been heavily attacked by contextualists. The common stance has been that the use of language is much richer and much less systematic than formal semantics can account for. The semanticist response, since Grice's theory of implicatures, has been to separate a core semantic interpretation, which delivers *what is said*, from pragmatic processes that take place after semantic interpretation is in place.

This strategy has in turn been criticized as inadequate, since pragmatics enters in various ways already before implicatures, but not restricted to ordinary indexicality. These pragmatic effects have been called *modulations*, including free enrichment, loosening, and semantic transfer. They are not semantically controlled, but on the contextualist picture they enter into the determination of primary linguistic content, what is said.

Compositional semantics still has a role if there is a more restricted level of content that is compositionally determined, and from which the pragmatically modulated level of what is said can be derived. Such a level has been identified as *the minimal proposition*, where only values of indexical context dependence are added to literal meaning. However, François Recanati has recently defended his claim that the minimal proposition plays no role in an account of linguistic communication: composition of contents does take place, but at a pragmatic level.

In this paper I provide independent reasons for the existence of a different level of compositionally determined semantic content, a conceptual structure, which is needed whether or not Recanati is right about the minimal proposition. Given the conceptual structure as the core semantic output, pragmatic modulation enters at a stage when semantic composition is already performed.

1 Introduction

Peter reads in the morning newspaper

- (1) Yielding to pressure from President Bush, the European Union has put off plans to lift its arms embargo on China this spring.¹

As a result Peter comes to entertain a number of thoughts, including the thoughts *that the European Union has put off plans to lift an arms embargo*.

¹New York Times web edition 23 Mars 2005.

(1) was written by Steven R. Weisman. Let's assume, which is overwhelmingly plausible, that Mr Weisman entertained very similar thoughts, and that his thinking these thoughts partly explains why he wrote the sentence for the newspaper. The convergence of content between speaker, Mr Weisman, and hearer, Peter, constitutes communicative success. There was success, despite the fact (let's assume) that prior to these events, neither Mr Weisman, nor Peter, had seen the sentence, nor entertained any of the thoughts involved. Let's call this phenomenon *novel communicative success*. By common sense standards of success, novel communicative success is very common. We would like to know why, given the low prior probability of such frequent convergence.

A proposed explanation, dating back at least to Frege, appeals to semantic structure. Frege writes

It is astonishing what language can do. With a few syllables it can express an incalculable number of thoughts, so that even a thought grasped by a terrestrial being for the very first time can be put into a form of words which will be understood by somebody to whom the thought is entirely new. This would be impossible, were we not able to distinguish parts in the thought corresponding to the parts of a sentence, so that the structure of the sentence serves as an image of the structure of the thought. (Frege 1923, opening paragraph)

According to Frege, on the most natural interpretation, novel communicative success is made possible because of an *isomorphism* between the sentence

and the thought it expresses. That is, both the sentence and the thought are structured entities, and there is a 1-1 correlation R' between constituents t of the sentence and constituents m of the thought, and between structural relations σ of the sentence and structural relations η of the thought, such that a structural relation holds between two constituents of the sentence if, and only if, a corresponding structural relation holds between their R' images in the thought. That is, we have the Frege isomorphism principle

(F) If $R'(t_1, m_1)$, $R'(t_2, m_2)$ and $R'(\sigma, \eta)$, then $\sigma(t_1, t_2)$ iff $\eta(m_1, m_2)$.

This provides the beginning of an explanation provided certain conditions are met. Here I shall focus on the pre-establishment requirement. There must be a *pre-established* 1-1 correlation R between sentence constituents and thought constituents, such that either R' is identical with R , or else R' can be straightforwardly derived from R in the context of the event. If there is a pre-established correlation (and some other conditions are also met), then it is possible for speaker and hearer to be know, or otherwise be equipped with, only the *elementary* part of R , i.e. the restriction of R to elementary sentence constituents (simple expressions) and elementary thought constituents (concepts), and elementary structural relations, such as the syntactic subject-predicate construction and semantic predication (or simple satisfaction). The complex elements of the correlation can then be *constructed* from the elementary ones when needed. This pre-established elementary restriction of R is *shared* between speaker and hearer, and they

are able to communicate successfully by building up more complex structures in a coordinated fashion.

By contrast, if there is no such pre-established elementary part of R , then we don't understand how the existence of a correlation will help in the first place, for we don't know how speaker and hearer manage to achieve the *same* syntactic-semantic correlation. It will be no help for achieving communicative success if the speaker has her correlation R_s and the hearer his correlation R_h , with not much overlap between them. If there is a shared pre-established elementary part, then we have at least the beginning of an understanding of how it can be achieved. The prior establishment of the pre-established part is to be accounted for, if possible, by language learning theory or the theory of radical interpretation.

The pre-established correlation R is the *linguistic meaning*-relation, i.e. a standing, context independent relation between linguistic expressions (or terms) and (usually) extra-linguistic entities.² That this is also how Frege viewed the matter is a plausible assumption, since Frege regarded a thought, *ein Gedanke*, as both the content of a mental thought-occurrence and as the linguistic meaning of a sentence, provided the sentence does not contain indexical expressions.

Frege's suggestion in the quoted passage pointed out the way for formal semantics, the study of syntactic-semantic structures, to contribute to the theory of linguistic communication. In this respect formal semantics has,

²On the present view, this is just what linguistic meaning is explicated as.

however, met with much resistance from the contextualist point of view. In general, contextualists, from Wittgenstein and Austin onwards, have claimed that ordinary language use is much richer, much more diverse, much less structured and orderly than what is pictured from the perspective of formal semantics. Contextualist criticisms have taken more or less radical forms. In this paper I would like to focus on one particular contextualist critique, one recently defended by François Recanati in *Literal Meaning* (Recanati 2004).

Recanati's criticism can be fairly precisely stated in the present context: Recanati would agree that there is a shared pre-established syntactic-semantic correlation R , as well as a contextually constructed syntactic-semantic correlation R' , shared between speaker and hearer. However, Recanati denies that the R' -construction *derives* from an R -construction. On Recanati's picture, the constructions actually derived by speaker and hearer build on contextually *modified* semantic parts, and does not depend on any R -construction from the prior original parts. It seems then that the *complex* parts of the R correlation do not have any explanatory role to play. How the R' construction is derived is up to pragmatics to account for, without help from formal semantics.

Before turning to the details of Recanati's criticism, it will be necessary to set out with more precision the outline of the semantic explanation.

2 Compositional explanation

If there is a 1-1 correlation R between sentence constituents and thought constituents, for any sentence of a particular language L , then L has a *compositional semantics*. That is, it holds of L that

(**PC_i**) The meaning of a complex expression is determined by the meanings of its (immediate) parts and its mode of composition

This is an informal formulation of *the principle of compositionality*. Formally, the principle is often set out as

(**PC_f**) $\mu(\sigma(t_1, \dots, t_n)) = \rho(\sigma, \mu(t_1), \dots, \mu(t_n))$

Here μ is the semantic function, mapping grammatical terms on meanings. The t_i are grammatical terms, which are disambiguated expressions with structure made explicit, built up from simple expression and syntactic operator symbols, σ, σ_i etc. ρ is called a composition function, mapping a syntactic operator, or mode of composition, and the meanings of constituent terms on the meaning of the whole term. Since the Fregean correlation R' is 1-1, the syntactic objects in (F) must be unambiguous as well. (F) then entails (PC): if there is an isomorphism, then there is a composition function.

They are not, however, equivalent, for (PC) only states the existence of a *homomorphism* from terms to meanings, and this does not entail the existence of an isomorphism. (PC) allows two terms with different structure and different parts to be mapped on the same meaning, and in such a case there

is no isomorphism. Correspondingly, there can be a homomorphism from terms to meanings even if meanings themselves are not *structured* entities. Because of the homomorphism requirement, there are functions $\rho(\sigma)$ ³ from n -tuples of meanings to meanings, and hence a *meaning algebra*, but there is no requirement that the value of such a function have its arguments as constituents (just as the number 12 need not have 5 and 7 as constituents despite being their sum).

With the existence of an isomorphism we need structured entities in the meaning algebra as well as in syntax. From the perspective of explaining communicative success there is good reason for this. Compositionality can help explaining how the hearer manages to derive the content of a complex expression from the contents of the parts, but it does not explain how the *speaker*, who starts out with a content, manages to derive the suitable syntactic expression for dressing it in words. To explain this analogously (see [work of author 1]) we need something like an inversely compositional principle:

(**IPC_i**) The expression of a complex content is determined by the expressions of the parts of the content and its mode of composition.

Such an inversely compositional principle requires that the idea of a part of a content makes good sense, and this requires structured contents. I shall not here set out the inverse principle formally, since one must then take

³That is, you get such a meaning function by applying the composition function ρ to a syntactic operation σ .

account of the complications of synonymies. The existence of synonymous expressions in the language means that the expression of a content cannot be uniquely determined, but only up to a certain congruence relation, which needs to be specified (see Pagin 2003a).

Since we need structured meanings as the immediate input for the speaker, we will also need them as the immediate output for the hearer. If the original semantic function μ does not deliver structured meanings, we can still achieve this by deriving from μ a meaning function μ' such that for simple expressions e $\mu'(e) = \mu(e)$ and for complex expressions it holds that

$$\mu'(\sigma(t_1, \dots, t_n)) = \langle \rho(\sigma), \mu'(t_1), \dots, \mu'(t_n) \rangle$$

We then get the original meanings back from the structured meanings by an evaluation function E such that where $\mu'(t)$ is an $n + 1$ -tuple r , and $p(i, r)$ is a function giving the i th member of r as value,

$$E(r) = p(0, r)(E(p(1, r)), \dots, E(p(n, r)))$$

If μ delivers propositions as the meanings of sentences, then we get propositions back from the conceptual structures that are the outputs of μ' by means of the evaluation function E . This idea will be put to use in the final section.

To have a workable theory, we would also need to account for context dependence. We must then distinguish between two levels of meaning, a standing, context independent, meaning, and the meaning or semantic value

an expression has in a particular context, its occasion meaning (as it is called in Pagin 2005). This is well known from David Kaplan's distinction between *character* and *content*. The occasion meaning of a complex expression, like the sentence

(2) I am eating

can be arrived at either by operating on the occasion meanings of 'I' (the speaker) and 'am eating' (true of eaters at the time of utterance), which are given from the context of utterance together with the standing meanings of these constituents, or by deriving the standing meaning of the sentence itself, and getting its occasion meaning from its standing meaning and the context. The difference is not of great significance. If there is nothing more to context dependence than assigning occasion meaning to context sensitive expressions, then the original Fregean correlation R relates expressions to standing meanings, while the derived relation R' relates expressions and occasion meanings. As we shall see, however, there is more. I shall not here elaborate the semantic account of this restricted kind of context dependence (see Pagin 2005).

The combination of compositionality and inverse compositionality has been called *strong* compositionality. It is reasonable to interpret Frege as having intended something like strong compositionality in his suggestion for an account of novel communicative success.

Strong compositionality cannot, however, be directly applied in an explanation. The problem is that the mere appeal to abstract meanings cannot

explain anything about speakers' performances, which are physical and mental. That (PC) holds of a language does not by itself say anything about how hearers manage to attach the right content to expressions. The abstract meanings must be connected somehow with what speakers do or think.⁴

I have proposed (Pagin 2003a,b) that the gap be bridged in the following way. A working explanation of novel communicative success is *possible* if syntax and semantics are so related, according to our syntactic theory **T** and our semantic theory **M**, that we can add a theory of mental representation **G**, such that

- i) **G** systematically assigns a representation of each element in the domains of **T** and **M** respectively, and
- ii) **G** systematically assigns a mental operation to each part of the definition in **M** of functions between the domains of **T** and **M**.

G provides mental operations mapping representations of syntactic elements on representations of semantic elements (and vice versa). If our theories of syntax and semantics are such that a theory **G** can be added which meets these conditions, then we do have an account of how actual communicative success *can* be achieved. And for this we need only consider syntax and semantics themselves.

⁴It has become fairly common in recent years to think of contents themselves, like concepts, as mental entities. This is fine in itself, as long as one does not hope one term to denote, without ambiguity, both mental and abstract entities. If we are dealing with communication, then there must be a content that is *shared* between speaker and hearer. Speakers don't share mental entities or events. If contents are what is communicated, they are not mental.

Putting these ideas together, we get the picture of a hearer that from a given utterance of a sentence s compositionally builds up a mental representation of a compositional structure that is the immediate occasion meaning of s , and then perhaps moves on to evaluate this structure into a proposition, or representation of an external state of affairs, such as the state of affairs *that yielding to pressure from President Bush, the European Union has put off plans to lift its arms embargo on China this spring.*

I shall now turn to Recanati's criticism.

3 Recanati's rejection of Minimalism

In his recent book, Recanati directs several attacks at truth conditional semantics, but I shall here focus on the composition objections. As always, the strategy of contextualist attacks on formal or truth conditional semantics is that of providing examples of utterances to show that literal or truth conditional meaning helps less than formal semantics needs to explain how the hearer understands the speaker. When literal meaning is not enough, pragmatic phenomena take up the slack. In extreme versions, literal meaning does not help at all (or does not exist).

An utterance of

(3) It is raining

can illustrate pragmatic processes of two different kinds, which Recanati calls 'saturation' and 'modulation' respectively. The first kind, saturation, is

the process of providing values to parameters in linguistically controlled and mandated way (Recanati 2004: 7-10). It is mandated because it is necessary in order to get a truth evaluable proposition. For instance, the sentence (3) does not by itself express a proposition that can be evaluated as true or false. We at least need to add a time, which is normally provided as the time of utterance. This is an example of saturation. So described, this is the process of arriving at what was called ‘occasion meaning’ in section 2.

It is usually held that a location needs to be provided as well, so as to express a proposition of the form

(4) It is raining at time t at place l .

If a location value is required, then adding one is also an example of saturation. If a location value is not required, and the speaker does not provide one, then by an utterance of (3) at a time t she expresses that it is raining *somewhere* at t , which is true or false. If a location value is optional, as Recanati holds, then providing a location value is an example of what Recanati calls ‘free enrichment’ (Recanati 2004: 10). Free enrichment is a kind of *modulation*. Modulations in general go beyond saturation. They are not required for reaching a truth evaluable proposition, but are pragmatically added.

Another kind of modulation is *loosening* (Recanati 2004: 24), as in

(5) The ATM swallowd my credit card

where the verb ‘to swallow’ has its application conditions extended to include

the cash machine process referred to.

A third kind of modulation is *semantic transfer*, exemplified by

(6) The ham sandwich left without paying

as said by one waiter to another at a restaurant (the example is due to Geoffrey Nunberg). Here the phrase 'the ham sandwich' is used to refer to the guest who ordered the ham sandwich, rather than to the dish itself. The semantic value has been transferred from the latter to the former.

Modulation is characterized as a top-down process, where the hearer uses his general understanding of the situation to arrive at the interpretation, as opposed to a bottom-up process, where the additions are triggered by elements of the sentence used.

Saturation and modulation are what Recanati calls *primary* pragmatic processes. *Secondary* pragmatic processes are distinguished by taking the result of the primary process as input. The secondary processes comprise various kinds of *implicature*. An example of Recanati's combining all three kinds is

(7) I've had breakfast

said in reply to the question 'Do you want something to eat?' (Recanati 2004: 8). The speaker of (7) *implicates* the she is not hungry and hence does not want anything to eat. Thereby it exemplifies implicature. It also exemplifies free enrichment, since she communicates not only that she has had breakfast at some time or other prior to the time of utterance, which

by itself is already true or false, but over and above this that she has had breakfast at the very day of the utterance.

According to the position Recanati calls ‘minimalism’ there is a real communicative level of *what is said* that minimally deviates from the literal meaning of the sentence used. The minimal deviation comes from saturation.⁵ Recanati rejects minimalism, arguing that the only real level of *what is said* is that delivered by the *combination* of saturation and modulation (Recanati 2004: 21).

More crucially, the rejection of minimalism is backed by the claim that the so-called *minimal proposition*, the result of saturation, is *not*, or at least *not* in general, computed by the hearer. In the example (6), according to Recanati, the hearer does not first interpret the utterance as an assertion of the absurd proposition that the dish itself left without paying, as minimalism implies, but achieves the semantic transfer locally, i.e. pragmatically reinterprets the phrase ‘the ham sandwich’ before the combination with the interpretation of the verb phrase is performed (Recanati 2004: 30-31). He later comments on the consequences for formal semantics:

Contrary to what formal semanticists tend to assume, the (intuitive) truth-conditions of our utterances are not compositionally determined by the meanings of words and their semantic arrangement, in a strict bottom-up manner. They are shaped by contextual expectations and

⁵There is further view called ‘syncretism’ by Recanati, according to which there are two levels of what is said, the minimal one and a further level that results from modulations. Recanati rejects this position as well (2004: 64).

world-knowledge to a very large extent. That is true of all utterances, however ‘literal’ they are (in the ordinary sense) (Recanati 2004: 81-82).

In the remaining two sections, I shall discuss whether Recanati’s claims are correct, and what their consequences are for semantics.

4 Always modulation?

Before turning to the consequences of modulation as part speaker-hearer performance, let’s briefly consider the question whether or not modulation always takes place. It seems not, for it is hard to see what enrichment, semantic transfer or loosening takes place in a straightforward utterance of

- (8) At five o’clock AM on Friday 25 Mars 2005, President Chirac of France was asleep.

Of course, you can go on to claim that we really have a number of tacit enrichments, since the speaker probably meant that Chirac was asleep in some normal way, in a bed standing in a house rather than under some highway bridge in the Alps, not under sedation, not chained to the radiator by a housekeeper, and so on (along the patterns in the contextualist literature), but in fact the speaker need not have tried to communicate any such thing. Even if the speaker in fact *believed* that Chirac was asleep in a normal bed, in a normal house, this need not be something she was trying to convey.

As far as I understand, Recanati does not insist that modulation always takes place.⁶ Recanati does object in other ways to traditional ideas of literal meaning, for instance by appeal to Searle’s notion of *background*, or Waissman’s notion of *open texture*, as applied to lexical meaning (chapter 9), but these are considerations of a very different nature. It may well be that we would not know how to apply the verb ‘to sleep’ to a genetically changed being whose brain hemispheres alternated in being in sleep mode. But if anything, this is then a restriction in the semantics of the verb, not an indication that tacit enrichment is taking place in the course of communicating with it.

I think we can safely conclude that there are null-modulation utterances, the interpretation of which can then apparently be explained by a theory that includes a compositional semantics. It may well be, on the other hand, that such utterances are very rare. In that case, if we want to explain the high *rate* of novel communicative success, the occasional occurrence of a null-modulated utterance is negligible. For instance, sentence (1), an ordinary news report formulation, contains such things as ‘pressure from President Bush’, which probably should be treated as applying loosening to the meaning of ‘pressure’, since it has nothing to with physical forces, as well as to

⁶There are hedges in some passages indicating this, like in ‘Natural language sentences, by contrast, express a complete thought (say something definite) only with respect to a context of utterance—in many cases at least.’ (Recanati 2004: 83); ‘[...] even if the semantic value of a word is fixed by language (and context, if saturation is necessary), composing it with the semantic values of other words *often* requires help from above. It is semantic composition which has a fundamentally pragmatic character’ (Recanati 2004: 139), my emphases.

the meaning of ‘put off’, since it is applied to an organization rather than to an animate individual, and so on. So it is necessary to consider the cases where modulation does take place.

5 Composition and modulation

In general, a pragmatic modulation can operate on a whole sentence content as well as on parts. We have an example of the latter with the semantic transfer on ‘the ham sandwich’ in (6). Another standard example, due to Kent Bach, illustrates an enrichment operating on the content of the whole. A mother says to a child who has just suffered a minor cut

(9) You are not going to die

conveying that the child is not going to die *because of that cut*. In Recanati’s terms, this is an enrichment of the contextually saturated meaning of (9).

We can then represent the outcome of modulation and composition in the same format as (PC). I shall use ‘ $C(t, c)$ ’ as referring to the modulated content of the term t in context c , where modulation might be null-modulation. I shall use π_i as modulating operators on contents. Then we can represent the combined operation as

$$(\mathbf{PMC}_f) \quad C(\sigma(t_1, \dots, t_n), c) = \pi_i(\rho(\sigma, C(t_1, c), \dots, C(t_n, c)))$$

saying that the modulated content of the term $\sigma(t_1, \dots, t_n)$ in context c is a function of the modulation operation π_i on the value of ρ as applied to σ

and the contextually modulated contents of the constituent terms. At the lowest syntactic level we have modulations directly of the meanings of simple expressions.

Since it is not determined from syntax and simple features of the context which operations the π_i are, but requires pragmatic competence of speaker and hearer, there is a question whether the purely semantic structure is of any help. The question is then whether, on the hearer side, (PMC) simply *replaces* (PC), in cases of modulation, or whether (PC) is needed in the account

5.1 Pragmatic composition and semantics

Assume that composition really takes place at the pragmatic level, according to (PMC). Does semantic composition still have an explanatory role? What role we assign to semantics will depend on whether or not we think that semantic *composition* has a role to play in selecting the modulation functions. But even in the extreme case, where modulation is completely local, the answer is affirmative.

It is affirmative because formal properties of the pragmatic composition are determined by the semantics. First, in all the modulation examples given, the composition function, by which simpler contents are mapped on more complex contents, is simply taken over from the semantics. If the speaker starts out with a subject-predicate construction, semantically interpreted as a concept applying a to an object, the modulated combination will again cor-

respond to the subject-predicate construction, with some concept – perhaps a different one – applying to some object – again perhaps a different one.

Secondly, and connected with the first point, in the standard modulation examples presented, the semantic *category* of the original meaning is preserved under modulation. That is, if the original meaning is a concept of an action, the modulated content will be a concept of an action. If the original meaning is a sortal concept, then the modulated concept will be a sortal concept as well, and so on.⁷

Third, the structural role of a modulated part in determining the modulated proposition will be the same as the structural role of the unmodulated part in determining the minimal proposition. The enriched interpretation of a grammatical direct object will again be the interpretation of the grammatical direct object.

So, structural features of the pragmatic interpretation are determined by literal semantic structure. Since modulation also, as Recanati thinks, depends on the literal meanings one starts out with, semantic structure will explain much of pragmatic interpretation. Semantics does not just offer a bag of literal meanings.

⁷There are rare exceptions to this, as when you say after Oscar has been accidentally squashed (‘Oscar’ was the proper name given to a tomato) that “There was Oscar all over the floor”. In that case the concept of an object is mapped on the concept of a stuff, but in this case too, the composition function is taken from the linguistic construction, and the non-category preserving modulation of an object concept is required by the semantic composition.

5.2 Modulations after semantic composition

The importance of semantic structure for pragmatic interpretation speaks in favor of the interpretation model proposed in section 2: the hearer first generates a conceptual structure determined by syntax, literal meanings, and perhaps contextual values for indexicals. Here I would like to consider a possible objection to compositional semantics that is made irrelevant by this model.

On Recanati's way of distinguishing between primary and secondary pragmatic processes, only the latter are *inferential*. With implicatures you reason from a propositional input to a propositional output. It might be thought that an inferential model could not work for modulation, because the nature of the input does not allow operations on *constituents*. The idea is that once the saturation has been carried out, what we have reached is just an integral, unstructured proposition. In case the interpretation is absurd, the hearer infers that the speaker meant something else, and has to reason his way to it. But the reasoning principles don't provide any operations on *parts* of the occasion meaning. For example, the semantic transfer of the reference of 'the ham sandwich' is not available as an operation on the (absurd) proposition *that the ham sandwich left the restaurant without paying*. Therefore, the argument concludes, the modulation must take place at the construction stage, not as an inferential process afterwards.

However, this is irrelevant if the immediate semantic output is a conceptual structure. In that case, each element of the structure is accessible for

modulation, and it is possible to define operations on the whole structure that effect modifications of its parts.

As an extra bonus this approach gives us a natural account of our sense of basic understanding of semantically or pragmatically anomalous sentences, like Searle's

(10) Cut the sun

which Recanati also appeals to. We do understand it in the sense of compositionally deriving a conceptual structure, like $\langle \text{apply}, \langle \text{apply}, \text{the, sun} \rangle, \text{cut} \rangle$, without perhaps knowing how to evaluate this structure, in sensible way, into a representation of a state of affairs.

5.3 Is modulation strictly local?

On Recanati's account, during the primary pragmatic process the hearer associates contents with expressions, and does not derive them by inference. The content associated with a particular expression is the one that in the context is most *accessible* (Recanati 2004: 30-34). Sometimes, the literal (saturated) meaning is most accessible, and sometimes an enriched or loosened alternative is more accessible and will be what the hearer associates with the expression. In the case of (6), the interpretation *the ham sandwich orderer* is more accessible to the hearer than the literal interpretation. But it is natural to object, and Kent Bach has objected (Bach 1994: 158), that the hearer does seem to process the absurd minimal proposition, for why would

he otherwise prefer the modulated interpretation over the literal? Bach adds that if the modulated interpretation is more accessible to begin with, and processed locally, then it would be more accessible also for a sentence like

(11) The ham sandwich is getting eaten

with the result that the sentence is absurdly interpreted as meaning *that the ham sandwich orderer is getting eaten*.

Recanati replies (Recanati 2004: 33-34) that modulation *is* sensitive to linguistic context. The predicate of the sentence, ‘left’, or ‘left without paying’, requires a person denoting subject term, and noticing that, the interpretation *the ham sandwich orderer* becomes more accessible to the hearer than the literal interpretation. But, according to Recanati, the hearer need not *process* the minimal proposition for this change to take place. It is enough to notice the linguistic context.

However, this does not seem right. For the linguistic context can be expanded so as to reverse accessibility again:

(12) The ham sandwich left without being payed for

In this case the literal interpretation of ‘the ham sandwich’ is right, despite the verb ‘left’ as head of the predicate. Rather, ‘left’ here undergoes loosening so as to take non-animate arguments, but this does not become clear until the modifier ‘without being paid for’ has been interpreted. Moreover, to complicate things further, this is clear if the orderer was alone, but if he

wasn't, and his company stayed behind, then in (12) again 'the ham sandwich' might be taken as referring to the orderer as well as the argument of 'being payed for', in case the interpretation of this phrase undergoes loosening to the effect of including persons *for whom* something is paid. All in all, it seems to be a complex interaction of linguistic elements with linguistic *and* non-linguistic context. The choice of interpretation of a single constituent does not in general seem to be determinable by noting the presence or absence of a simple categorial mismatch.

This elaboration suggests that whether or not to effect modulation, and what modulation to effect, is a much more global affair in the process of interpretation than Recanati in crucial passages admits. In other passages (Recanati 2004: 30-32) he acknowledges e.g. that there might be reason to back-track and reverse an earlier interpretation, but his only comment is that this is compatible with an accessibility model: accessibility orderings may change as more material is interpreted. However, this is irrelevant to the question of whether the minimal proposition gets processed or not. If interpretation is global, then before a final decision on modulation is made, the entire sentence, and in some cases maybe even larger stretches of discourse, must be interpreted.

These considerations do not give much evidence for what normally does happen during interpretation, and finding out does not seem easy, but no good reason is given for the view that the hearer does not derive the minimal proposition.

But whether or not he does, this is not decisive for the role of compositional semantics. On the present model, the hearer in the first stage, compositionally, generates a conceptual structure from the entire sentence or another large syntactic unit, and then in the second stage, in the course of *evaluating* this structure into a flat propositional representation, effects various modulations on elements of the structure so as to reach the most reasonable interpretation of the speaker in the given context. On this model, modulation is not part of the primary composition of contents. On this model, again, compositional semantics plays a vital role as an interpretational core engine in the communication process.

Affiliation

References

- Bach, K., 1994, ‘Conversational implicature’, *Mind & Language* 9:124–62.
- Frege, G., 1923, ‘Gedankengefüge’, *Beiträge zur Philosophie des Deutschen Idealismus*. 3:36–51. Reprinted in Frege, *Logische Untersuchungen*, Vandenhoeck & Ruprecht, Göttingen, 1976. Translated by R. Stoothoff as ‘Compound thoughts’, published in *Mind* 72 (1963), 1–17.
- Pagin, P., 2003a, ‘Communication and strong compositionality’, *Journal of Philosophical Logic* 32:287–322.

Pagin, P., 2003b, 'Schiffer on communication', *Facta Philosophica* 5:25–48.

Pagin, P., 2005, 'Compositionality and context', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*, 303–48, Oxford University Press, Oxford.

Recanati, F., 2004, *Literal Meaning*, Cambridge University Press, Cambridge.