

Overhearing a sentence

Recanati and the cognitive view of language*

Fernando Martínez-Manrique and Agustín Vicente

Universidad de Granada / Universidad de Valladolid

Many pragmaticians have distinguished three levels of meaning involved in the comprehension of utterances, and there is an ongoing debate about how to characterize the intermediate level. Recanati has called it the level of ‘what is said’ and has opposed the idea that it can be determined semantically — a position that he labels ‘pragmatic minimalism’. To this end he has offered two chief arguments: semantic underdeterminacy and the Availability Principle. This paper exposes a tension between both arguments, relating this discussion with Carruthers’s cognitive view of language, according to which some thoughts are, literally, sentences of our natural language. First we explain how this view entails minimalism, and we construct an argument based on semantic underdeterminacy that shows that natural language sentences do not have the compositional properties required to constitute thoughts. Then we analyze the example of a subject’s overhearing a sentence without an interpretive context, arguing that in the light of the Availability Principle the corresponding thought can be regarded as a natural language sentence. Thus, semantic underdeterminacy and availability pull in different directions, and we claim that there is no characterization of the latter that can relieve this tension. We contend that Recanati’s availability shares with Carruthers’s introspectivism an overreliance on intuitions about what appears consciously in one’s mind. We conclude, therefore, that the Availability Principle ought to be abandoned.

1. Introduction

In trying to distinguish different levels of meaning involved in the processing of utterances, many pragmaticians have hit upon a threefold distinction. Roughly characterized, the distinction places at one extreme the meaning of the sentence (or sentence-type) itself, and at the other the speaker’s meaning,

or the meaning that a particular speaker intends to convey by means of a particular utterance of a sentence on a particular occasion. In between, we find an intermediate level that is also directed to utterances yet tries to capture in them a relatively general, stable meaning.

One controversial account of this intermediate level has been set forth by Recanati (2001, 2003), who employs the expression ‘what is said’ to refer to it, as opposed both to ‘sentence meaning’ and ‘what is implicated’, which are his terms for the other two levels. In order to characterize the intermediate level, a critical issue will be how much context is needed to characterize its meaning. According to Recanati, responses cluster around two poles, one that tries to keep context to a minimum, another that exploits context maximally, as a necessary component to reach what a speaker says. Recanati has labeled these approaches, respectively, *minimalism* and *maximalism*, and has launched a sustained attack against the former (Recanati 1995, 2001, 2002a, 2002b, 2003).

His criticisms cluster around two notions: *semantic underdeterminacy* and the *Availability Principle*. The argument from semantic underdeterminacy says, in a nutshell, that what is said in most utterances is underdetermined by the semantic meaning of the uttered sentence. On the other hand, the Availability Principle holds that ‘what is said’ “must be analyzed in conformity to the intuitions shared by those who fully understand the utterance — typically the speaker and the hearer, in a normal conversational setting” (2001: 80). The meaning delivered by these intuitions will be generally richer than semantic meaning. As minimalism relies heavily on semantics, it does not have the resources necessary to capture what is said.

As Carston (2002b: 132) points out, Recanati’s notion of ‘what is said’ denotes very much the same sort of entity as Bach’s ‘implicature’ (Bach 1994) or Sperber and Wilson’s ‘explicature’ (Sperber and Wilson 1986/95). It is also close to the ‘utterance meaning’ that Dascal (1983) adapts from Grice. But there are significant differences in their respective treatments. For Bach implicatures go beyond what is said, even if they are built out of what is said (Bach 2001: 19). Sperber and Wilson propose explicatures as explicitly communicated assumptions that consist in combinations of linguistically encoded and contextually inferred conceptual features (Sperber and Wilson 1986: 182), while Recanati contends that ‘what is said’ is obtained non-inferentially. Dascal’s utterance meaning is dealt with by semantics (more specifically, by semantic pragmatics) and corresponds to what a speaker *conventionally* meant by an utterance, including conventional implicatures, so it is not necessarily identical with what is said — understood as what is explicitly said (Dascal 1983: 36–37), while for

Recanati 'what is said' is pragmatic through and through, and does not include implicatures.

The debate is loaded with a number of different issues, such as how to set the boundary between pragmatics and semantics, or whether those levels are reflected in actual stages of utterance processing that take place during interpretation. In this paper we want to take a different path and relate Recanati's arguments with a current proposal by Peter Carruthers (1996, 1998, 2002) regarding the involvement of natural language in thought.

According to Carruthers, we make use of natural language (henceforth NL) as a *vehicle* of some of our thoughts. That is, some thoughts that we have are, literally, sentences of a NL. This position has been labeled the *cognitive view of language*.¹ The basic argument Carruthers (1996) offers for this view is that our own introspection reveals that we do use language when we think². Put in a very simplified form, what Carruthers says is that (a) introspective data reveal that sometimes we think in a NL and (b) we must begin by taking these data for what they seem. Moreover, Carruthers argues that we should go further than what introspection reveals — that is, that some episodic conscious thoughts are linguistically codified — and admit that we also use NL to codify latent thoughts and unconscious token-thoughts of the same types as those episodic thoughts.

There are two points that allow us to connect this discussion with the minimalist/maximalist debate in pragmatics. First, Carruthers's cognitive view of language embraces a minimalist account of what is said, indeed, a radical, literalist one. Therefore, arguments against minimalism can be employed to undermine his conclusions. We will argue that this is exactly the case, offering an argument based on considerations of semantic underdeterminacy that shows that NL cannot be the vehicle of thought. Second, Carruthers's introspectivism takes seriously people's intuitions about what appears consciously in their minds when they are having a particular thought. In a parallel way, Recanati's Availability Principle takes seriously people's intuitions about the content they are consciously aware of when they understand a particular utterance. Both theories give primacy to something that is located at the *personal* level, rather than at the *subpersonal* one. Consequently, success in both theories depends on having the right picture of the elements of thought at the personal level.

The way we will develop these points is the following. Sections 2–4 develop the connection between the minimalist character of Carruthers's cognitive view and the considerations against minimalism based on semantic underdeterminacy. In Section 2 we will review Recanati's chief arguments

against minimalism, focusing on the two aspects already mentioned: semantic underdeterminacy and the Availability Principle. In Section 3 we will present Carruthers's proposal in more detail and see how his cognitive view entails minimalism. Section 4 will offer an argument that goes from the semantic underdeterminacy of NL to its non-compositionality, and from this to the inability of NL to constitute thought.

The remaining sections of the paper develop the connection between the Availability Principle and the nature of whatever elements are consciously brought to mind when one entertains a given thought. Section 5 analyzes the example of a subject's overhearing a sentence without an interpretive context, and raises the question of what thought, if any, the subject can be attributed to in such a case. We will argue that it is perfectly compatible with the Availability Principle that the thought entertained be regarded simply as a natural language sentence, just as the cognitive view of language claims. But if this reading is allowed, then we reach conclusions that pull in the opposite direction from those obtained when we apply the argument of semantic underdeterminacy. That is, semantic underdeterminacy leads to the conclusion that the interpreted content of an utterance goes far beyond the semantics of the sentence. In contrast, the case of overhearing offers a putative example in which the hearer's intuitions about the content of the heard utterance are limited to the sentence itself, and forces one to reconsider whether the sentence is semantically underdetermined after all. So we can perceive a tension between the conditions established by semantic underdeterminacy and the Availability Principle. Section 6 examines this tension, arguing that Recanati lacks a proper characterization of availability that allows him to avoid the damaging conclusions against semantic underdeterminacy. Finally, Section 7 deals with two possible objections to our analysis of overhearing — its non-generalizability and its separation of meaning from comprehension — and in trying to solve them brings light on further considerations with respect to the Availability Principle and why it should be rejected.

2. Against minimalism

Recanati characterizes pragmatic *minimalism* as the doctrine according to which “what is said’ departs from the meaning of the sentence (and incorporates contextual elements) *only when the sentence itself sets up a slot to be contextually filled*” (Recanati 2001: 77). On this account, what is said corresponds

to a minimal proposition that is achieved by semantic means. To be sure, minimalism grants that in order to obtain the minimal proposition contextual elements are often needed, but they belong to a limited class.

Following Bach (1997: 39), Recanati distinguishes between two sorts of context: “*wide context* concerns any contextual information relevant to determining the speaker’s intention (...); *narrow context* concerns information specifically relevant to determining the semantic values of indexicals and is limited to a short list of contextual parameters. Narrow context is semantic, wide context pragmatic” (Recanati 2002a: 111).³ The minimal proposition that captures what is said needs to resort to the narrow context to assign values to the slots set by the sentence. Yet this type of context-dependence is semantically tractable: “insofar as the contextual assignment of values to indexicals proceeds according to linguistic rules (...), there is no reason not to consider that aspect of content-determination as part of *semantic* interpretation” (Recanati 2002a: 110). Recanati employs the term *saturation* for the process by which the meaning of a sentence is completed and made propositional. In the minimalist picture, other contextual factors, those that belong to the wide context, are dealt with only after the process of saturation has taken place and has yielded a minimal, truth-valuable proposition. In other words, wide contextual elements will be needed only to reach the speaker’s meaning — that, in Recanati’s terms, corresponds to what is implicated.

The view opposed to minimalism is pragmatic *maximalism*. It claims that ‘what is said’ includes contextual elements that go beyond the slots set by the sentence, i.e., it is much richer than the minimal proposition allowed by the minimalist. There are two kinds of pragmatic processes at work: *primary* non-inferential processes that help determine what is said, and *secondary* inferential processes that take ‘what is said’ as input and deliver ‘what is implicated’ as output. Among primary processes some are mandatory, and others are optional. Saturation is a paramount example of the former, while free enrichment, a process that allows the addition of extra elements according to a particular context, is the most typical optional one.

The important thing about optional pragmatic processes is that they do not come into play *after* a semantic interpretation has been obtained (say, by means of saturation). Rather, they interact with semantic interpretation in an inseparable way: “we cannot separate those aspects of speaker’s meaning which fill gaps in the representation associated with the sentence as a result of purely semantic interpretation, and those aspects of speaker’s meaning which are optional and enrich or otherwise modify the representation in question. They are

indissociable, mutually dependent aspects of a single process of pragmatic interpretation” (Recanati 2001: 88).

To see the contrast between minimalism and maximalism, consider the utterance:

(1) I’ve had breakfast

What the utterance says in the minimal sense is that the speaker has had breakfast at some point in the past. Yet the meaning that the utterance usually conveys (e.g., as an answer to ‘Are you hungry?’) is that the speaker has had breakfast this morning. For the minimalist the extra element [this morning] does not belong to what is said: it is part of the speaker’s meaning, typically involving an inferential process that takes the minimal proposition (what is said) as the starting point, and an enriched proposition (what is implicated) as the inferred meaning.

In contrast, maximalism contends that elements like [this morning] must be regarded as a rightful part of what is said, not as a product of implicature. The reason is that the content that we try to capture at this intermediate level corresponds to the *intuitive* truth-conditions of the utterance, i.e., “to the content of the statement as the participants in the conversation themselves would gloss it” (Recanati 2001: 79–80). Hence, Recanati claims that a proper account of ‘what is said’ must observe his *Availability Principle*, which holds that ‘what is said’ “must be analyzed in conformity to the intuitions shared by those who fully understand the utterance — typically the speaker and the hearer, in a normal conversational setting” (2001: 80).

Recanati rejects a possible compromise between both positions (what he calls the *syncretic view*) by means of which the intermediate level they are trying to capture is split in two: a minimal proposition and, subsequently, a maximally enriched one. The reason for this rejection is that, as we pointed out, for Recanati what is said is obtained by indissociable processes: there is no psychological room for any minimal or literal meaning that the interpreter must decode at an early stage. His attack against minimalism is mainly based on two broad kinds of arguments: semantic underdeterminacy, and the availability condition.

2.1 Semantic underdeterminacy

Recanati’s contention is that the phenomenon of semantic underdeterminacy cannot be treated in the way suggested by minimalism.⁴ This phenomenon

refers to cases in which a sentence contains an expression whose meaning is underdetermined until we introduce pragmatic factors from the wide context in its interpretation (see also Bach 1994; Carston 2002b; Sperber and Wilson 1986/95). This is especially conspicuous in the case of demonstratives, like ‘that book’, but many other expressions follow the same pattern. For instance, the expression ‘John’s car’ does not have a determinate value until we do not take into account what sort of relation between John and the car the speaker has in mind. Whether it is a relation of ‘ownership’ or ‘*drivership*’ or ‘being the car John bet on’ or any one of numerous other possibilities, is something that can only be determined by looking at the wide context. This includes any element of the interaction between speaker and interpreter (e.g., perceptually salient elements, previously uttered sentences, their common personal histories, and so on) that can make a particular semantic value more accessible than any other candidate. In other words, we cannot turn the relation of John to his car into a semantic rule, because there is no rule that can tell us which of the possible elements will be most accessible on each occasion: “What a given occurrence of the phrase ‘John’s car’ means ultimately depends upon what the speaker who utters it means. It therefore depends upon the wide context” (Recanati 2001: 85).

One might wish to say that there can be something like a ‘default rule’ for the relation that ‘John’s car’ expresses.⁵ Maybe it is possible to pinpoint one of the possibilities, say, ownership, as the relation that is routinely assigned to expressions like that at an initial stage. This could constitute the ‘literal interpretation’ of the relation expressed by [’s]. After this stage, and when the context reveals that the interpretation of the compound was inappropriate, one may backtrack and select a better choice.

However, Recanati shows that there are cases in which this strategy does not work. He asks us to consider the example ‘the lion’s sword’, where ‘the lion’ refers to a particular person (a warrior with a lion painted on his shield) that used the sword (in contrast with an alternative interpretation in which ‘the lion’ refers to a particular animal that was killed with the sword). In a case like this, the alleged literal interpretation (say, [’s] interpreted in terms of ownership of the sword) cannot be determined until it has been resolved what is nonliterally meant by the speaker (i.e., that ‘the lion’ refers metonymically to a particular warrior). The utterance, Recanati contends (1995: 229), is “semantically indeterminate at the purely literal level”, i.e., there is no literal level to begin with until we give a nonliteral value to some constituent. For if we adopt the ‘default semantic rule’ strategy in this case, i.e., assigning default literal values to all constituents, the result is something like ‘the sword owned by the lion (the

animal)’. This is, in Recanati’s terms, a monster that cannot be plausibly entertained at any stage of processing. The minimal proposition, he concludes, is an abstraction with no psychological reality.

At this point, the defender of minimalism may agree that there are special cases that demand a more elaborate processing: a processing in which, as Recanati advocates, elements from the wide context are taken into account. But in most instances of normal communication, the minimalist can insist, semantic interpretation does deliver a psychologically kosher interpretation that provides a suitable starting point for inferential processes to work on. Yet, Recanati contends that *most* cases, even in normal communication, follow the same pattern as ‘John’s car’ or ‘the lion’s sword’. Semantically underdetermined expressions appear in more sentences than one might initially suspect, as the following paragraph suggests:

“Now I take it that such [semantically underdetermined] expressions can be found all over the place. Moreover, semantic underdetermination is not limited to particular lexical items. One can follow Waismann and argue that the satisfaction conditions of any empirical predicate are semantically underdetermined and subject to pragmatic interpretation. There is also constructional underdetermination. For example, consider something as simple as the Adjective + Noun construction, as in ‘red pen’. What counts as a red pen? A pen that is red. But when does a pen count as red? That depends upon the wide context” (2002a: 112).

If Recanati’s generalization is right, then minimal propositions are the exception rather than the rule. So any realistic account of language processing ought not to rely on them.

2.2 The Availability Principle

The minimal proposition cannot constitute ‘what is said’ by the speaker because what is said has to meet what Recanati calls the ‘Availability Principle’: “what is said *must* be available — it must be open to public view (...). Hence, (...) ‘what is said’ must be analyzed in conformity to the intuitions shared by those who fully understand the utterance — typically the speaker and the hearer, in a normal conversational setting” (Recanati 2001: 80). To put it another way, what matters to assess the truth or falsity of what is stated in a particular utterance are the intuitive truth-conditions as perceived by those that have access to the context in which the sentence was uttered. The intuitive truth-conditions are likely to differ from the minimal truth-conditions: even if the sentence is not

semantically underdeterminate, i.e., even if a minimal truth-evaluable proposition can be obtained by purely semantic means, it will be often the case that its truth-conditions are not what we need to take into account to determine the truth or falsity of the utterance. Consider again example (1):

(1) I've had breakfast

If it is the case that the speaker did not have breakfast *that morning*, in normal circumstances we will conclude that he is not telling the truth, even though he may have had breakfast three weeks ago. What is implicated in the two cases will also differ. If (1) is an answer to 'Are you hungry?', I am entitled to understand that the speaker is not hungry when I interpret his utterance in terms of the intuitive truth-conditions, but not when I interpret it in terms of minimal truth-conditions.

The Availability Principle has been criticized on a number of counts. Bach (2001) holds that even if intuitions provide data for semantics, it is unclear whether they reveal semantic facts or whether they play a role in ordinary communication. García-Carpintero (2001) offers a dispositional account of availability that allows minimalism to comply with this principle: minimal propositions would be dispositionally available, inasmuch as competent language users can reflectively abstract the systematic contributions of linguistic units in different contexts of utterance. Carston (2003) claims that the Availability Principle applies to pragmatic processes generally and is thus unable to sustain Recanati's distinction between 'what is said' and 'what is implicated'.

In our view, even if one accepts the principle's main tenet — that only intuitive truth-conditions matter to determine what is said — the notion of 'intuitive truth-conditions' requires subsequent clarification. There are two sides to it, public and private. The public side has to do with the demand that what is said must be "open to public view". It corresponds to "the content of the statement as the participants in the conversation themselves would gloss it." (Recanati 2001: 79–80). In our opinion, a way to read this is that what is said corresponds to the common answer that each and all of the participants would give, when asked, precisely, "what did the speaker say?". It is interesting to note that in this task the speaker counts as one more vote. He does not have the last word, even if what he meant was something different from what the rest of the participants understood. If the context made a particular meaning *M* more likely, then it can be concluded that the speaker said *M*, even if this was not his intention.⁶

The private side of availability refers to the fact that, ultimately, what is said has to be consciously available to each individual: “the participants in the conversational process are aware of what is said” (Recanati 2001: 80). This makes void any attempt to safeguard minimal propositions by arguing that they are unconsciously entertained at some point in processing; this would grant them some psychological reality at the cost of rendering them useless to account for what is said. As Carston (2002b) points out, what is said in Recanati’s sense is determined at the personal level — the level at which thoughts are accessible to the person herself — rather than at the subpersonal — the structures responsible for thought and behavior that lie beyond the subject’s conscious grasp.

A point to note is that the Availability Principle and the thesis of semantic underdeterminacy are independent of each other. Semantic underdeterminacy without the Availability Principle is possible: most sentences may not have determinate semantic truth-conditions, yet the truth-conditions that matter may be minimalist, and not the intuitive ones. Conversely, we might hold the Availability Principle without having semantic underdeterminacy: every utterance may be purely semantically determinate, yet what is said may correspond to intuitive truth-conditions that go beyond semantic meaning. Hence abandoning one principle does not entail abandoning the other. This is relevant for our purposes because we think that the Availability Principle should be discarded, and the thesis of semantic underdeterminacy retained. In the next two sections we will review the cognitive view of language and we will argue that it is undermined by an argument from semantic underdeterminacy. In Section 5, in contrast, we will offer a reading of the Availability Principle that helps to vindicate the cognitive view of language and hence questions that previous undermining. We will then give further reasons for rejecting that principle.

3. The cognitive view of language

The view that we think in our own public language is not new. Explicit versions of it can be traced back to Hobbes and Condillac, and it plays a prominent role in the well-known Sapir-Whorf thesis. More recent articulations of the view can be found in Sellars (1969), Harman (1975), Dummett (1989) or Gauker (1994), to name a few. Yet for at least three decades the view has been unpopular in mainstream cognitive science, which was dominated by the idea that thought has its own specific medium — a universal medium that possibly constitutes a language but that is distinct from, and expressed by, the different

public languages. There is, however, an increasing number of authors that defend the involvement, in different degrees, of natural language in cognition. To offer a quick selection, we have approaches from pragmatics that regard language not only as an instrument of communication but also as an instrument of thought (Dascal 1983, 2002);⁷ linguistic and psycholinguistic approaches that show the influence of language in domains such as spatial cognition (Lucy 1992; Levinson 1996); experimental data from psychology (Spelke and Tsivkin 2001a, 2001b) and interpretations of the archeological record (Mithen 1995) that suggest that our mother tongue may be acting as an intermodular *lingua franca*; evolutionary arguments (Bickerton 1995); considerations of mental architecture (Dennett 1991) that support a role of language for certain types of conscious thought; and broadly Vygotskian philosophical reconstructions that propose that we use natural language as a way to approach our thoughts in a reflexive way (Clark 1998). Our focus in this paper will be limited to the views of Peter Carruthers (1996, 1998, 2002), as someone who has defended a strong position about this topic.

3.1 Thoughts as natural language sentences

According to Carruthers, we make use of natural language⁸ as a *vehicle* of some of our thoughts. That is, some thoughts that we have are, literally, sentences of a NL. This position has been labeled the *cognitive view of language*. The basic argument Carruthers (1996) offers for this view is that our own introspection reveals that we do use language when we think. Carruthers offers a weak and a strong version of his thesis. The weak version says that “some human conscious thinking (*viz.* conscious propositional thinking) is such that, of natural necessity, it involves public language, in virtue of the given architecture of human cognition together with causal laws” (1996: 252). The strong thesis adds that “necessarily, some of these propositional thoughts belong to types which (...) constitutively involve such language” (1996: 263–264), i.e., that there are thoughts that could not have been tokened in the absence of language. Even though Carruthers offers some arguments for the strong thesis,⁹ it is only the weak version that he regards as solidly established. Still, the weak version has strong consequences and, as we will see soon, it is a species of the literalist view.

First, to dispel a possible misunderstanding about the term ‘constitutively’, it must be noted that *both* the strong and weak thesis take it that some thoughts are constituted by public language in the sense that they are *tokened* in NL sentences (and not in a distinct language of thought). The sense in which the

strong thesis is strong is by adding that linguistic tokening is necessary for those thoughts to be tokened at all. Second, for Carruthers many thoughts are conducted *only* in public language sentences. This is a consequence of his introspectivist thesis, according to which some thoughts appear in consciousness purely as sentences. Third, even those thoughts that do not appear consciously as linguistic have often a linguistic component. The obvious instances are imagistic thoughts, where ‘image’ is understood in a broad sense to cover visual, auditory or spatial experiences, among others. Carruthers (1996: 36) accepts that some thoughts can consist entirely in those images. But he claims that there is one kind of thinking, *propositional thinking*, which cannot be conducted in purely imagistic thought. Propositional thinking always involves a public language component, and in many cases it is carried out completely by linguistic means.

The fact that Carruthers endorses the claim that many thoughts are constituted entirely by NL sentences may be obscured by the space he devotes to the possibility of mixed representational vehicles, i.e., vehicles that combine language and images. In fact, this discussion is part of his effort to argue that many imagistic thoughts are language-dependent. This dependence is manifest in a double sense. First, many images need to be embedded in a linguistic context in order to have a determinate content (1996: 253). This is due to the fact that images are very restricted about the content they are able to convey. Second, some of them are thoughts *about* images, and these thoughts depend on language to be entertained (1996: 254). Carruthers lays great emphasis on his arguments that language is present in conscious imagistic thoughts, probably because he takes as self-evident that language is present in linguistic thoughts. What he is at pains to show is that even those thoughts that do not appear readily as linguistic rely often on a linguistic bedrock.

To sum up, even if one rejects the strong thesis that many thoughts are impossible to entertain without language, the weak thesis put forward by Carruthers holds that NL sentences are *de facto* the primary vehicles of our conscious thoughts. Mental sentences do not provide access to contents that are the “real components” of thought; rather, the sentences are constitutive of thoughts themselves. Otherwise one would never entertain thoughts in a mode in which one had immediate access to their contents, i.e., one would never have conscious thoughts at all. We do not wish to discuss the theory of consciousness on which these views are based. We will concentrate instead on the consequences about the nature of thoughts that Carruthers’s views entail. In particular, we will show its literalist lineage.

3.2 The cognitive view as a literalist thesis

Literalism holds that sentences have literal meanings.¹⁰ A standard way to spell out this idea is by means of meanings that result from the application of conventional rules of composition on the conventional meanings of their parts.¹¹ It is possible to construct a non-mentalistic reading of this thesis, so that literal meanings are regarded as linguistic abstractions that have nothing to do with the processing of language. On the other hand, a mentalistic version holds that literal meanings play some role in the mental processing of utterances. In other words, whatever is the “final” meaning extracted by the interpreter of an utterance, there is some stage at which the literal meaning is entertained.

The cognitive view of language (Carruthers 1996) is literalist in this sense. Its main tenet is that if there is introspective evidence that a person *seems* to be thinking in a NL sentence, then she *is* thinking in a NL sentence. Hence, if the sentence

(2) Bill crashed into the bank

is introspectively perceived in her mind, then we can take (2) *as a thought* of hers. In other words, she *is literally* thinking (2). Thus, if we grant the cognitive view its first premise, i.e., that we often seem to think in NL sentences,¹² then we have literal meanings all over the place, because every time one thinks in a NL sentence one is *ipso facto* thinking that sentence literally.

Now, it might be objected that a phenomenon as ordinary as ambiguity immediately challenges this literalist reading. For instance, when someone entertains (2), she will be thinking two very different things depending on whether the bank in question refers to ‘building hosting a financial institution’ or ‘the side of the river’. So her thought cannot consist *only* in (2). However, the cognitive view does not deny that there is another level, possibly an unconscious one, where the ambiguity is solved: the occurrent thought includes dispositions associated to the imaged sentence. But those dispositions may well involve relations — especially inferential relations — to other linguistic structures. The imaged sentence constitutes the core of the occurrent thought (1996: 63), and it makes no sense to say that “real thought” takes place at another place. So it is still the case that, at some level of the processing of meaning, the thinker entertains, as a thought, the literal sentence (2). This is all that a literalist reading needs.

As a second objection, it might be remarked that perhaps an ambiguous sentence like (2) is never actually entertained, at any level. So when a person utters (2) what he consciously has in mind (as opposed to what he utters) is something like, say,

- (2) b. Bill crashed into a river bank

The reason why he utters (2) would be that the context allows him to be more economical. Likewise, when a person hears (2) what she consciously understands (as opposed to what she hears) is (2b), provided that the context makes it clear that Bill was driving a motorboat. So while ambiguous sentences like (2) are often *uttered*, perhaps they may never be *entertained*. What one entertains is only a contextually appropriate rendition. But even granting this, the cognitive view is still a literalist thesis. This is so because those “appropriate renditions” that are entertained in mind are still NL sentences, and, as such, they are entertained literally. There is no use in trying to complicate matters arguing that the “appropriate rendition” is more complex than (2b): if it is a NL sentence, and if the subject consciously entertains it, then she is literally thinking that NL sentence.

Hence, the only way to counter the cognitive view is to show that the “appropriate rendition” is *not* a NL sentence. One way to do that is to demonstrate that NL sentences do not have the right properties that allow them to be constitutive of thought. We will do this by means of an argument based on considerations of semantic underdeterminacy.

4. From semantic underdeterminacy to non-compositionality

The apparent semantic underdeterminacy of natural languages offers a strong argument against Carruthers’s view. In a nutshell, the argument is this: (i) Cognition requires a semantically precise and compositional instrument or vehicle; (ii) but expressions of NLs are largely semantically underdetermined; and (iii) if an expression is semantically underdetermined, it is not compositional.¹³ So, from (i) and (iii), NLs cannot be the instrument of cognition.

In 3.1. we saw Recanati’s reason for (ii). With respect to (iii), the argument from semantic underdeterminacy to non-compositionality can be summarized as follows:

- a. (*Compositional Conditions*) Compositionality requires that the meaning of the whole is obtained solely from the semantic value of the parts, plus the rules of composition.
- b. (*Semantic Underdeterminacy*) There are cases in which the parts of a NL sentence (i.e., the lexical components) do not have a definite semantic

value until we add pragmatic factors. There are other cases in which the rule of composition is not definite, until we add pragmatic factors.

- c. (*Generalization*) Point (b) is generalizable to most NL sentences.

Hence

- d. Compositionality does not hold for most NL sentences.

Recanati's examples of demonstratives like 'that' in 'that dog', or syntactic relations like [s] in 'John's car', are aimed to demonstrate that NL is intrinsically ill-designed to provide a determinate content. In other words, they support premise (b) above: 'that' is an indexical whose value is indefinite without pragmatic factors, while [s] affords no definite rule that specifies how to compose the relation of John to the car. Inasmuch as these cases are pervasive in NL sentences (as premise (c) claims), then most of them do not meet the conditions to be compositional. If Recanati's examples are correct, then semantic underdeterminacy denies the compositionality of NL *tout court*.

There are some possible considerations that may tell against this argument. First, one may resort to the role of extra-linguistic elements in fixing the content of a mental sentence. The idea would be that a combination of public language with those elements is not subject to underdeterminacy. Carruthers might be ready to bargain for this weaker position, given that he phrases his thesis as "conscious thinking *involves* public language". This may suggest that even if language is unable by itself to fixate a determinate content, the underdeterminacy is dispelled once a sentence is filled in with extra-linguistic components. Yet, the only extra-linguistic elements that Carruthers considers are images. Images may be useful to complete a number of incomplete sentences like 'I want this chair to go there'. The idea is that by inserting an image first after 'this chair' and then after 'there', we may get a determinate, compositional sentence. However, many underdetermined sentences are not determinable in this way. For example 'John's car is empty' is doubtfully completable by means of an image, just as 'all the students have gone on strike', which needs a domain in order to be truth-evaluable. Moreover, as we said above, in Carruthers's view *it is language that lends determinate content to images*, rather than the other way round. His introspectivism leads him to the consequence that many conscious thoughts appear both as "purely linguistic" and as determinate. But the argument we have just presented denies that public language can be determinate in the required sense. So it seems that neither NL nor NL-plus-images can be the vehicle of thought.¹⁴

A second consideration is that it is possible to view literal meaning in a way that does not rely on compositionality. The idea is that literal NL sentences may be mentally entertained even though they are not compositional. Dascal (1987, 1989) proposed such an alternative view of literal meaning, which he called ‘moderate literalism’. He contends that “the idea of a *complete* determination of literal meaning by compositional means is unwarranted” (1987: 261), and offers instead a list of cumulative criteria, such as, context invariance, non-cancelability, conventionality, and compositionality itself. The more criteria are satisfied by an aspect of meaning, the more we can regard it as part of literal meaning. Furthermore, lack of compositionality does not prevent public language from being used in cognition.¹⁵ The point is, thus, that the cognitive view can escape the underdeterminacy argument by endorsing the claim that literal meaning plays a role in cognition despite its putative non-compositionality. However, the role that the cognitive view wants for NL is stronger than what Dascal’s moderate literalism grants. It is one thing to argue that language is *used* by cognition in some way or another — a claim that we are sympathetic with — and quite a different one to affirm that it is *constitutive* of thought. Carruthers’s position is committed to the latter view, and this is the view that the underdeterminacy argument challenges.

Still, to complete the argument against the cognitive view of language, we have to argue for (i), i.e., the idea that thought must be compositional. If compositionality is a requirement for thought, and if NL is not compositional, then NL cannot be constitutive of thought. Fodor (2001) has provided an argument on behalf of the compositionality of thought. A thought, he claims, is always explicit (or exact) about its content because “a thought *is* its content”. Unlike natural language, sentences in the language of thought do not have to be interpreted by a further court. Consequently, each Language-of-Thought sentence can have *only one content*.

Fodor (1998) conceives of thoughts as being composed out of conceptual atoms. These ultimate components of Mentalese, unlike NL words, are context-independent. There is no further court to whose interpretation they must defer. Hence, while having a word does not entail having its content, because we need a context to determine the latter, having a concept amounts to having its content, which is unique. But this is generalizable to whole thoughts: a thought also has a unique content that is determined by the contents of its component concepts. Consequently, we need rules of composition that allow the construction of thoughts from concepts, in a fixed way. These rules must be context-independent (i.e., the content of the compound $f(X, Y)$, where X

and Y are concepts, and f is a compositional function, is the same regardless of context). Therefore, having a thought is equivalent to having its content, while in order to obtain a (truth-conditional) proposition from a NL sentence we need something else, a context.

To sum up, whatever serves as a vehicle of thought cannot be inexplicit about the content of its sentences, so that a thought may bear exactly the content it is meant to bear. To this end, and given that thoughts have to be formed from a finite set of pieces, a vehicle of thought has to be compositional. As a NL cannot be compositional in this strict context-independent sense, it cannot be a vehicle of thought. In this way the semantic underdeterminacy argument threatens the minimalist view and, with it, the cognitive view of language. The argument also implies that stronger literalist views cannot succeed either. The cognitive view, we argued, is a literalist view of this kind; hence if minimalism is wrong, Carruthers's thesis must be wrong.

However, the cognitive view of language may meet the availability criterion. As we have already mentioned, what Carruthers's project has in common with Recanati's is the insistence on taking as the relevant data thoughts as they appear at the personal level, i.e., to subjects themselves. The sort of psychological reality that matters to what is said is conscious grasp. In the following section we will analyze a case in which what is consciously grasped is precisely the type of structure that Carruthers takes as constitutive of thought: a natural language sentence.

5. Overhearing a sentence

Suppose you are on a train and overhear a conversation in which someone utters sentence (3):¹⁶

(3) John's car is green

You do not have any further background to interpret this sentence: you do not know who John is, or what relation he bears to the car. What sort of thought do you entertain when you hear (3)? We think that there are three main possibilities:

- i. You do not have the context to determine what sort of relation [$'s$] is, so you do not reach any interpretation and do not entertain any thought whatsoever.

- ii. You contribute a context yourself, say, you interpret [’s] as a relation of ‘ownership’, so that you entertain a thought whose content is, roughly, that the car owned by John is green.
- iii. You do not contribute a context, and entertain a thought that corresponds to the NL sentence that you heard, namely, that John’s car is green.

We take it that (i) is counterintuitive: there is certainly *something* that you, as a competent English speaker, understood in hearing (3). Possibility (ii) looks more promising: you know nothing about John or his car, so you make up a plausible scenario. However, if one wants to reject minimalism, there is a tempting idea that one should leave aside: the idea that the context you contribute to the interpretation of the overheard sentence is a default context (something like “relation [’s] *typically* expresses ownership”). The reason, as we said in Section 2.1, is that default rules are convertible into semantic rules, which is all that minimalism needs in order to work. Default contexts, then, have to be considered part of the narrow context, not of the wide context, and the narrow context is semantically tractable.

So, a non-minimalist reading of (ii) has to say that when you overhear (3), the slot set by the possessive phrase can be filled by practically anything in the wide context, from environmentally salient elements to whatever the interpreter has in her imagination. Recanati’s own maximalist theory (Recanati 1995) proposes an *accessibility-based* sequence of utterance processing. In his model, the literal values of constituents are initially recovered by the hearer, but not a literal interpretation of the *whole*. Instead, the hearer can reach first a nonliteral interpretation of the utterance, provided that the elements of this interpretation are associatively derived, by spread of activation, from the literal constituents that were initially accessed. Consequently, the interpretation for which you actually settle will depend on which of those elements are more accessible to you.

There is a problem, however, for this reading. If context-selection is a matter of accessibility, and given that it *must* be the case that some contextual element is more accessible than the rest, then whatever element meets this condition provides the context. Then, if the overhearer provides herself with a scenario to interpret the overheard utterance, the accessible elements that matter in order to fix the context are whatever thoughts happen to be in the overhearer’s mind at the moment of overhearing. The outcome is that interpretation in this case becomes quite a random process. For instance, if it is the case that while on the train you have been thinking about your nephew’s latest painting, then you might interpret ‘John’s car’ as ‘the car painted by John,’ given that the relation

'the X painted by Y' was the latest relation that you used to fill relation [s]. We take it that an interpretation of this kind is very unlikely, i.e., it is psychologically implausible that the overhearer will entertain such wild thoughts.

Our claim, thus, is that, if the overhearing subject were asked to "guess" a plausible interpretation for (3), she would restrict her answer to a very limited range of choices. We conjecture that the answers of most subjects would converge on the same set, most likely, relations of 'ownership' and 'drivership'. In fact, as Dascal (1989: 255) points out, experiments show that in a "no-context" situation subjects tend to understand utterances in their literal interpretations. This raises the question, thus, of "how wide" the wide context is. If not everything that happens to be accessible in the wide context can be used to interpret an utterance, then the notion of accessibility has to be constrained somehow. However, it is difficult to see how this can be done without paving the way to minimalism. If we can substantially reduce the number of choices of the interpreter, then we can obtain something like a default rule for the constituent, i.e., we can make it semantically tractable. In other words, saying that it is possible to constrain the interpretations of [s] for which one may tolerably settle amounts to saying that there is something *intrinsic* to [s] that restricts its probable semantic value.

The maximalist has a possible reply. Recanati's model, we said, is accessibility-based. It is possible to conceive of a parallel model in which "all candidates [semantic values] that reach a certain level of accessibility are tried in parallel; the first candidate whose processing yields satisfactory results in the broader context of discourse is retained, while the others are suppressed" (Recanati 1995: 211).¹⁷ As the overhearer is in a situation of uncertainty, she may entertain several interpretations, none of which seems totally satisfactory. She does not settle for any particular interpretation, remaining uncertain about what the speaker really meant, i.e., the overhearer does not form any particular thought.¹⁸

However, we have no more reason to think that the overhearer does not settle for any of the interpretations, than we have to think that she settles for the first one that comes to her mind. According to Recanati (1995: 226), an interpretation is satisfactory, and hence retained by the interpreter, when it "fits the broader context of discourse" in a coherent way. As in the case of overhearing the context is almost non-existent, the outcome is not that *none* of the interpretations will be satisfactory, but that *any* of them will be. At any rate, it seems to us that there is an economical, plausible alternative: that the overhearer adopts a cautious, non-committal, stance, and settles for an interpretation

that corresponds just to the sentence meaning. So if we asked her ‘what did the passenger say?’, she would just *repeat the utterance* just as she heard it, e.g., she would answer ‘the passenger said that John’s car is green.’

This phonetically full-blown NL sentence meets the Availability Principle. With respect to its private side, the NL sentence is phenomenologically available to the overhearer’s consciousness. Here the evidence seems to support Carruthers’s introspectivism: the most immediate elements that we are aware of are strings of words.¹⁹ As for its public side, it is arguable that the most public entity is the uttered NL sentence itself. This is consistent with the common sense view that one obtains the most accurate rendition of what was said on a particular occasion if one is able to recall the *exact* words that were uttered.

The claim, therefore, is that there is a psychologically real structure available to the overhearer at the personal level, and that this structure corresponds, literally, to a NL sentence. This has consequences for the thesis of semantic underdeterminacy. This thesis states that it is not possible to obtain the meaning of an expression if we do not resort to pragmatic, contextual factors that are not semantically tractable. But the case of overhearing suggests that it is possible to entertain a thought corresponding to the meaning of an utterance even in the absence of context. Semantic underdeterminacy gave us grounds to conclude that NL does not have the properties necessary to constitute complete thoughts. Availability, on the other hand, allows us to hold that ‘what is said’ can be captured by NL sentences, inasmuch as these are available, and, given that what is said corresponds to a thought, NL sentences correspond to thoughts themselves. There seems to be a tension between the conditions established by availability and semantic underdeterminacy respectively.

6. A tension between availability and semantic underdeterminacy

Here is the problem again. Both the Availability Principle and the thesis of semantic underdeterminacy are concerned with ‘what is said’ by an utterance: the former states that what is said must be available to conscious and public view, the latter that what is said cannot be obtained by purely semantic means, even less correspond to a NL sentence. But given that NL sentences meet the availability conditions, they are plausible candidates for constituting what is said, contrary to what follows from semantic underdeterminacy. This tension suggests that one of the theses, availability or underdeterminacy, should be abandoned. We propose that it is the former that must go, but let us first

examine several ways — unsuccessful, in our opinion — in which ‘availability’ could be construed so that the tension might be resolved.

One may note that being available is a necessary, but not sufficient, condition to constitute what is said by a given sentence. This is the case, for instance, of meanings that are obtained from implicatures. According to Recanati, what is implicated is obtained inferentially from what is said, and it also has to be available to consciousness.²⁰ So not everything that is available belongs to what is said. However, one would not like to say that NL sentences, to the extent that they are available, belong to the level of what is implicated. They would belong to a third category of meaning: something that is available — and hence mentally entertained — yet different both from ‘what is said’ and ‘what is implicated’. Yet, for Recanati these two levels exhaust the types of “meanings in mind” admissible in his theory. There is only a third kind of meaning, sentence meaning, but this one has no psychological reality for him, being a mere linguistic construction. Now, if one accepts that NL sentences are consciously available, but rejects that they belong either to what is said or to what is implicated, what one is effectively doing is to reinstate sentence meaning in the psychological realm. But this move results in a vindication of literalism — the view that it is possible to obtain the literal meaning of a sentence — because it is difficult to see what “psychological realistic sentence meaning” could be, if not literal meaning. To sum up, a construal of ‘availability’ in terms of necessary but not sufficient conditions is not helpful to save this notion.

Another construal of ‘availability’ may hold that when we say that a NL sentence is available we are using a notion of phenomenological availability, and it may be objected that this is too strong a notion. The crucial issue for Recanati would be whether language users are aware of the truth-conditional content of the utterance, that is, whether they are able to specify the conditions under which the utterance is true.²¹ In terms of this less demanding notion, what appears phenomenologically in their minds has little bearing on what is said. A hearer is aware of what is said inasmuch as the truth-conditions of utterances are available to him. Applying this rationale to the overhearer, what matters is if she can specify the truth-conditions of the overheard sentence. As she lacks the appropriate context, she cannot specify those conditions. So even though she can quote the sentence literally, what is said by that sentence is not available to her.

However, it is disputable that the overhearer cannot specify the truth-conditions of the sentence. Consider the following specification for sentence (3):

‘John’s car is green’ is true if and only if John’s car is green

One may complain that what appears on the right side of the formula is underspecified. In particular, one may demand a specification of the relation *R* that holds between John and the car. However, demands for specification have to stop somewhere. For instance, it would not be reasonable to demand that the particular shade of green be specified. It is clear that there is an indefinite number of green shades that can satisfy the predicate so as to make (3) true, and that (3) is not committed to a specific one. In a similar vein, it can be argued that there is an indefinite number of relations between John and the car that can satisfy the genitive relation, and that (3) is not committed to a particular instance of them. The interpretation of a sentence like (3) could be roughly equivalent to ‘a certain car that bears a certain close relation to a certain subject called John is green.’ The interpreter is not committed to give further assignments to many of the elements of the sentence. This does not mean that she gives no assignment at all: the value assigned to the relation [’s] is very abstract indeed, since it amounts to something like ‘a genitive relation between John and a car’. This is expressed in a much more economical way by means of the words with which the relation was expressed in natural language, i.e., expressing it as ‘John’s car’. In fact, public language’s own underdeterminacy may be instrumental in preserving some unspecificity of the content we want to transmit (see note 15).

A third different way to construe the notion of availability may arise from paying attention to the fact that its public and private sides do not always seem to fit smoothly. Perhaps what is available is better captured by one of these sides. However, if we consider what is privately available sometimes we end with disparate intuitions. Given that different interpreters may have different mental contexts, they may be aware of different meanings. Consider utterance (4):

(4) Employees must wash hands

A typical hearer will understand it as a requirement for employees to wash *their own* hands. Arguably, ‘their own’ must be regarded as an element that belongs to the intuitive truth-conditions, and that is added by the process of free enrichment that we mentioned in Section 1. Yet we may conceive of a hearer, perhaps from a different cultural background, who understands the sentence as a statement about the duties of employees in a particular job, so that they are required to wash the hands *of customers*. If this is the meaning that is consciously accessible to the second hearer, then ‘of customers’ will also be part of the intuitive truth-conditions for this subject. We thus have two different private intuitive truth-conditions that correspond to two different conscious grasps.

Perhaps we can resort then to the public intuitive truth-conditions in order to apprehend what is said in such a case. However, there is no straightforward manner to characterize what is common to both subjects. If we assume that there is some “common meaning” grasped by all participants, the result may not be very different from the minimal proposition that the minimalist proposes. If we resort to some “average meaning” that stands somewhere between the richer and the poorer grasped meanings, we may end with a construct that has no psychological reality, contrary to what the Availability Principle tries to capture. If we allow that there are as many fillers for what is said as there are different interpretive contexts (i.e., that what is said may be different for each speaker/hearer dyad of the conversational participants), it is unclear whether the notion of ‘what is said’ is interestingly different from the notion of ‘what is implicated’. It is part of the distinctive character of the latter to allow for greater variation, given that a subject may obtain a wide range of implicatures from a particular utterance, but ‘what is said’ is supposed to remain more or less constant for each different context.

To sum up, availability seems to us too obscure a notion to serve as a criterion that helps to determine what a sentence says. If we characterize it in a strong way, including in it the phenomenological properties of whatever the hearer is aware of, then we are led to conclude that NL sentences are available in the required sense — i.e., that what is said is a NL sentence. This has the consequence of jeopardizing the thesis of semantic underdeterminacy, which we take to be a sound argument against the cognitive view of language. If we try an alternative, weaker characterization of availability, it is unclear that it can fulfill the task that Recanati intends for it. This task is to determine a conception of what is said that is, at the same time, richer than the minimalist conception, common to all conversational participants, and psychologically realistic. These three requirements seem difficult to maintain jointly. In the next section we will offer more reasons for rejecting the Availability Principle in the light of further objections that can be raised against our analysis of overhearing.

7. Further objections

To conclude this paper, we want to consider two objections to our analysis of overhearing cases. To be sure, overhearing situations are designed so as to prevent wide contextual elements from playing any role in the interpretation. But this is not the usual case, so one may wonder whether any lessons for normal

communication can be drawn from this example. On the other hand, our conclusions about what is entertained in overhearing cases seem to sever the link between the notion of ‘what is said’ and the notion of ‘understanding’. Let us treat these objections in turn.

7.1 Generalization

Overhearing cannot be regarded as a typical case of linguistic communication. In fact, it is not clear that it can be regarded as a case of communication at all, given that the overhearer cannot be included in the audience for which the speaker intended her utterance. Hence, whatever may happen in situations of overhearing cannot be taken as a model for the rest of communicative practices. There is normally a context available for interpretation, and the interpreter cannot help but rely on it. If the context, for instance, makes the relation of ‘betting on’ salient, the interpreter will automatically understand ‘John’s car’ as ‘the car that John bet on’, and not as the abstract, unspecific relation that the overhearer might entertain.

However, let us state once again what the overhearing example is intended to achieve. It simply asks us to take the place of a context-independent, non-committed interpreter of a sentence and to inquire “does this sentence have a psychologically real meaning for her?” Our claim is that the answer to this question is affirmative: a person typically entertains some thought when she overhears a sentence. The next step is to note that the mental structure that meets the Availability Principle is the NL sentence itself, as it consciously “sounds” in the overhearer’s mind. From these facts, one may conclude that the thought entertained by the overhearer corresponds just to the NL sentence. Now, it is obvious that in normal linguistic communication there are also NL sentences that are consciously entertained by the conversational participants. Given that those NL sentences meet the Availability Principle, what are the grounds for denying that they constitute the participants’ thoughts in the normal case too? This is the challenge that the overhearing example poses for Recanati’s availability-based account. And we contend that the way to deal with it is simply to reject the Availability Principle.

7.2 Semantics and understanding

Suppose that we accept that the meaning grasped in overhearing ‘John’s car is green’ can be characterized as the thought (framed in a NL sentence): ‘John’s car

is green'. This accounts for the fact that when we ask the overhearer what her fellow passenger said, she can answer: 'he said that John's car is green'. Suppose now that the sentence she overheard is

(5) John's snark is a bojum

where 'snark' and 'bojum' are not meaningful to her. Still, when we ask her what the passenger said, she can faithfully answer with the available NL sentence, namely, 'he said that John's snark is a bojum'. Going a step further, consider that the utterance was

(6) hobbledy gobbledy doo

It seems that the overhearer can still inform us that what is said in this assertion is 'that hobbledy gobbledy doo'.

Something must be wrong in an account that allows the interpreter to have a notion of what is said even in cases where there is no trace of understanding. We agree: something is wrong. We have tied the entertained proposition so closely to the NL sentence that in the end we seem to have identified 'what is said' with 'the string of sounds uttered', even if that string is empty of meaning. However, we claim that this wrongness is shared by an availability-based account of what is said. The source of the trouble lies in the assumption that identifies the psychologically interesting meaning — the content that is reached by the subject in the process of obtaining what is said by an utterance — with the mental structures that appear at the personal level — a level that is tied to a particular phenomenology. We suggest therefore that the assumption ought to be dropped. What is said by an utterance can be psychologically real yet be located at a subpersonal level. As Bach (2001) contends, conscious intuitions may provide data for semantics but they do not constitute the semantic facts themselves. To put it another way, just as the NL sentence that we see introspectively may be an expression of thought but not constitutive of it, the elements of which the hearer is aware when she interprets an utterance may be not constitutive of meaning either. Dascal's moderate literalism also provides an alternative way in which literal meaning may be psychologically real. The key idea is that literal meaning "need not be a *part of* the final interpretation of any given utterance" (Dascal 1987: 262) in order to play a role in the process of interpretation. Indeed, it is not necessary to see it as contributing content to the speaker's meaning. Rather, its role consists in "leading the hearer to the identification of the relevant items of contextual information which have to be used in order to come up with an interpretation" (*ibid.*). As we said above,

literal meaning may play this role even if it is not compositional. We can add now that it does not have to be consciously available either — indeed this will be the usual case.

Perhaps one can still see a fundamental difference in the overhearing example by reconsidering the intuition, expressed in Section 5, that there is something one understands when one overhears ‘John’s car is green.’ That is, perhaps possibility (i) in that section was correct: the overhearer does not entertain any thought. In Recanati’s terms, “what is wrong is the assumption that semantic interpretation can deliver something as determinate as a complete proposition” (2001: 85). In other words, maybe there is no corresponding thought to be entertained until we can assign more definite values to the constituents of the sentence.

As we already said, we find this idea counterintuitive, but some observations will be helpful to support our intuitions. First, let us note that Recanati seems to share our intuitions. Recanati (1997) argues that we can believe what we do not understand, and that the notion of belief involved is the same as ordinary belief. If this is so, then there must be some thought that corresponds to the non-understood belief. What is involved is a *deferential representation*, i.e., a case in which one cannot determine the truth-conditions of her belief but acknowledges that it has truth-conditions and that they are known by somebody else, to whose judgment she defers.²² Moreover, Recanati (1997: 93) contends that “deferential representations, though semantically determinate, are *epistemically* indeterminate: one does not know which proposition the deferential sentence she accepts expresses”.

In our view, the same treatment can be applied to overheard sentences. The overhearer apprehends a complete proposition and a complete thought, even if there are many ways in which the world can realize what is stated by the proposition. In fact, this “multiple realizability” would appear even if we tried to capture ‘what is said’ in maximalist terms — i.e., including in the interpretation elements from the wide context. For consider a contextually filled rendition for ‘John’s car’, say, ‘the car driven by John’. There is an indefinite number of ways in which this expression may refer to reality, i.e., many ways in which the car may be driven.

To see the issue from a different angle, let us make a distinction between *misunderstanding* and *non-understanding*. Misunderstanding means that the interpreter reaches a proposition whose truth-conditions are not those that the speaker intended. For instance, suppose that the speaker’s genuine audience (i.e., not the overhearer) understands ‘John’s car’ as ‘the car owned by John’,

while the speaker intended it as ‘the car driven by John.’ This is a case in which the interpreter forms a complete but incorrect proposition. Non-understanding, in contrast, means that the interpreter did not reach a complete proposition, and hence did not grasp a meaning at all. The clearest example of non-understanding is given by “sentences” like (6) that sound like mere gabbling. The claim, thus, is that overhearing is more similar to misunderstanding than to non-understanding. There is something that the overhearer understood, even if she cannot identify the truth-conditions intended by the speaker.

Perhaps one can insist that the case of overhearing is one of non-understanding, alleging that something like ‘a certain car that bears a certain relation to a certain subject called John is green,’ which we offered as an approximate rendition of what the overhearer understood, is too imprecise to be a proposition. One might insist that there is no thought that corresponds to it, only a nebulous schematic quasi-thought. Yet, most abstract thoughts share the same kind of impreciseness. In fact, suppose that one utters, in normal conversation, sentence (7):

(7) There is a certain car that bears a certain relation to John and is green

If we denied that there is a proposition that the speaker is trying to communicate, and that this proposition can be grasped by the audience, we would be denying the possibility of expressing and interpreting really abstract thoughts. In the case of overhearing, the overhearer may also reach an abstract thought — so abstract indeed that its truth-conditions may be satisfied in multiple manners.²³ We claim that the meaning that constitutes what is said by an utterance in a normal case is also an abstract object of this kind.

8. Conclusions

We have been arguing that it is possible to maintain the thesis of semantic underdeterminacy while rejecting Recanati’s Availability Principle. In our view, the argument from semantic underdeterminacy is really effective to undermine the cognitive view of language. Carruthers’s introspectivist argument seems to us to be mistaken because it equivocates between the vehicle of our introspected thought and the thoughts themselves. In other words, from the fact that the most public, externally available elements in conversation are strings of words does not follow that the intended contents are constituted by those strings of words. A sentence expresses a proposition but it is not a proposition. Similarly,

from the fact that the most conscious, internally available elements in mind are strings of words it does not follow that mental contents are constituted by those strings of words. A conscious “mental sentence” expresses a thought but it is not a thought.

Yet, Recanati’s Availability Principle suffers from a comparable equivocation. It conflates a notion of ‘what is said’ that corresponds to the utterance’s meaning with a notion of ‘what is said’ that corresponds to what is available to consciousness. But the meaning that the hearer reaches in interpreting an utterance does not necessarily correspond to the elements that she is aware of. There may be layers of meaning that have to be processed at a subpersonal level to reach a particular interpretation of an utterance. It is at one of those levels, we suggest, that what is said can be found. What the hearer entertains at the personal level may be too rich — too altered by contextual factors — to capture what the sentence says that is common to all interpretations. This does not necessarily commit us to the claim that ‘what is said’ will be ultimately couched in a purely semantic way, as the minimalist contends. Yet it is clear that it makes the job easier for a minimalist account in which, when a sentence is processed, a layer of meaning is devoted to a minimal proposition that is put to work for different purposes, namely, for all sorts of unconscious inferences. On the other hand, Recanati’s account faces a different kind of difficulty. Maximalism strives to offer an account of ‘what is said’ that is more complete than what minimalism has to offer, including in it the total pragmatic context. But there must be a point at which this process of completion stops. This point cannot be located at the highest maximal meaning that corresponds to the *most* specific content that the speaker has in mind, because this is a meaning that cannot be shared by conversational participants. So we need constraints for the context that can effectively enter the interpretation. But talking of constraints, as we have remarked several times, is minimalist-friendly, because it opens the door to default contexts, i.e., to tractable semantic accounts of the so-called semantically underdetermined expressions. It shows that there may be something in the semantics of words themselves that elicits the appropriate contextual factors to be taken into account.

Notes

* This paper is thoroughly collaborative. Order of authorship is arbitrary. This work is funded by Research Project BFF2002–03842 from MCyT and VA020/04 from the Junta de Castilla y León. We want to thank Marcelo Dascal and two anonymous referees of *P&C* for

their helpful comments, as well as audiences in the Universities of Granada and Valladolid where earlier versions of this paper were presented.

1. In contrast with the *communicative view* of language, which regards language merely as a means to express thoughts, not as a vehicle that constitutes them.
2. In Carruthers (2002) the argument from introspection appears only as the background for his position, while modularity-based considerations play the leading role. We will leave these considerations aside for the purposes of this paper.
3. A similar distinction between two types of contribution of the context was already made by Dascal (1983: 35ff): “that [context] which is involved in determining utterance meaning (which is a semantic role of the context), and that which is involved in determining the speaker’s meaning (which is its properly pragmatic role)”. He mentions Wunderlich and Searle as other sources of the distinction.
4. Recanati uses the alternative terms ‘indeterminacy’ and ‘underdetermination’ in his (2001) and (2002a), respectively. We will follow the usage of his (2003) throughout. See also Carston (2002a: 20) for a distinction between ‘indeterminacy’ and ‘underdeterminacy’.
5. Bach (2001: 39) pursues a related strategy, arguing that most uses of possessive phrases like ‘John’s car’ are derived from a “multifarious relation of possession” that has a certain primacy over the most unusual uses.
6. The speaker may still have a decisive vote, in cases when two or more equally suitable interpretations are available (as if the audience demanded: “What did you really mean?”).
7. Dascal introduced the term ‘psychopragmatics’ for the subdiscipline that has as its domain the study of the use of language in thought (Dascal 1983: 45ff).
8. Carruthers (1996: 261–263) makes a distinction between public and natural language. Since this distinction is irrelevant for our purposes, we will use the terms interchangeably.
9. His arguments are chiefly two: one based on the existence of thought constituents that depend on language for their acquisition, and the other based on considerations of the alleged greater simplicity and unifying power of his thesis.
10. Recanati (2002b) distinguishes between several types of literal meanings. These distinctions have no bearing on the view developed in this paper.
11. In Section 4 we will summarize a different approach to literal meaning (Dascal 1987, 1989), which is not biased to a compositional criterion.
12. This premise, however, is not uncontroversial. It can be objected that, in many cases, introspection does not reveal NL sentences, but fragments of sentences. Fragments do not have truth-conditions, so they do not seem to be fit as pieces of thought.
13. This point is something that follows from the notion of compositionality that interests us for the purpose of this paper. For a language has a compositional semantics in the present sense if the meaning of its sentences is a function of its parts, and the meaning of such parts

is not sensitive to wide contextual information: that is, if you can obtain the meaning of wholes just by looking at the meaning of their parts and their structure, while being blind to the wide context. We take it that the demand that the language of thought is compositional in this sense coincides (as we are about to show) with the Fregean demand that thoughts be “complete” or “determinate” in the sense of being truth-evaluable. We thank a referee for drawing our attention to this point.

14. In what follows, we will leave the role of images aside, since we think that it is a question that does not add much to the discussion. At any rate, we think that there is a second possible argument against Carruthers’s position, which works both against the use of NL and of NL-plus-images as vehicles of thought. If the major kind of evidence for defending that NL is the vehicle of thought is introspective data, and if the language of thought must be compositional, then a close look at what introspection reveals should convince us that we do not think in NL (or in NL-plus-images). We scarcely ever hear a whole sentence in our minds. We do not hold 19th century monologues, but utterly fragmentary Joycean monologues, where a single word stands for a whole thought. We do not have the space to develop this here (we have done so in Vicente and Martínez-Manrique (in press)) but we take it that these introspective data point to a different use of NL in thought.

15. Dascal (2002) distinguishes three ways in which language is involved in cognition: as an environment or context of use, as a resource that thinking can make use of, and as a tool that is engineered for specific cognitive tasks. Indeed, language’s indeterminacy can be an advantage when it comes to conceptualize vague domains or intuitions.

16. After writing this, we were reminded of the parallelism between our overhearing example and Katz’s ‘anonymous letter’ case (Katz 1977). However, we take it that they are devised with different purposes in mind, and that the respective analyses focus on different issues. Katz’s letter is mainly offered as a criterion to demarcate semantics and pragmatics, while our overhearing case addresses our intuitions with respect to the kind of thought that is provoked by an utterance in the absence of context.

17. Recanati also posits an accessibility-based *serial* model, in which instead of processing all the candidates simultaneously, the most accessible one is processed first. This difference is not significant for the point we are making.

18. Dascal (1983, 1989) has also suggested that processing of the three levels of meaning (in his terms: sentence, utterance, and speaker meanings) runs in parallel until the system settles for a satisfactory speaker’s meaning. In his view, this is compatible with the psychological actuality of the first two meanings: they are less accessible once the utterance was interpreted, but they were employed in achieving the interpretation, and they may be retrievable afterwards, if certain conditions apply.

19. Moreover, there are those who maintain (notably Jackendoff 1987, 1997) that conceptual structures are never themselves available to consciousness. Rather, they are expressed in conscious linguistic strings with a distinctive order and phonetic sequence. Consciousness, on this account, occupies an intermediate level between perceptual periphery and conceptual center. The same idea applies to other systems of representation, such as vision.

20. For Recanati, what is implicated has to be available in a stronger sense than what is said. The latter is the product of processes that work at the subpersonal level themselves, while the processes that produce what is implicated act at the personal level. This means that “the interpreter has to be aware of what is said, aware of what is implied, and aware of the inferential connection between them” (Recanati 2002: 114).

21. We owe this objection to an anonymous referee.

22. Religious beliefs concerning the “mysteries of faith” provide a particularly perspicuous case (see Dascal 1975).

23. An anonymous referee pointed out another possibility. Perhaps the overhearer entertains a thought, but the sentence, without a context, is semantically underdeterminate and does not express a thought. The thought which the overhearer entertains may result from existentially quantifying all the free variables which the sentence meaning contains (and which make the sentence underdeterminate). So from the fact that the overhearer, who hears the sentence without knowing the context, entertains a complete thought, it does not follow that the sentence, all by itself, expresses a complete thought. Thus the thought which the overhearer entertains might be different from the thought which the utterance expresses (what it says). We think that if this point is correct it helps to weaken Recanati’s position, which makes what the sentence says dependent on what the hearer entertains. Notice that the Availability Principle is precisely directed toward the hearer. It is possible thus, as we suggested, to maintain the thesis of semantic underdeterminacy while rejecting the Availability Principle. Bach (2001: 29) expresses a related concern: “how could it even seem that psychological or epistemological considerations about the *hearer* are relevant to the question of what the *speaker* says?”

References

- Bach, K. 1994. “Conversational implicature”. *Mind and Language* 9: 124–162.
- Bach, K. 1997. “The semantics-pragmatics distinction: What it is and why it matters”. *Linguistische Berichte* 8: 33–50.
- Bach, K. 2001. “You Don’t Say?”. *Synthese* 128: 15–44.
- Bickerton, D. 1995. *Language and Human Behavior*. Seattle, WA: University of Washington Press.
- Carruthers, P. 1996. *Language, Thought, and Consciousness: An Essay in Philosophical Psychology*. Cambridge: Cambridge University Press.
- Carruthers, P. 1998. “Thinking in language? Evolution and a modularist possibility”. In P. Carruthers and J. Boucher (eds), 94–119.
- Carruthers, P. 2002. “The cognitive functions of language”. *Behavioral and Brain Sciences* 25(6): 657–726.
- Carruthers, P. and Boucher, J. (eds). 1998. *Language and Thought: Interdisciplinary Themes*. Cambridge: Cambridge University Press.
- Carston, R. 2002a. *Thoughts and Utterances*. London: Blackwell.

- Carston, R. 2002b. "Linguistic meaning, communicated meaning and cognitive pragmatics". *Mind and Language* 17: 127–148.
- Carston, R. 2003. "Conversational implicatures and pragmatic mechanisms". Paper delivered at the meeting of the *European Society for Philosophy and Psychology, ESPP'03*.
- Clark, A. 1998. "Magic words: How language augments human computation". In P. Carruthers and J. Boucher (eds), 162–183.
- Dascal, M. 1975. "La razón y los misterios de la fe según Leibniz". *Revista Latinoamericana de Filosofía* 1(3): 193–226. [Reprinted as "Reason and the mysteries of faith according to Leibniz", in M. Dascal, *Leibniz. Language, Signs, and Thought*. Amsterdam: John Benjamins, 1987.
- Dascal, M. 1983. *Pragmatics and the Philosophy of Mind*. Amsterdam: John Benjamins.
- Dascal, M. 1987. "Defending literal meaning". *Cognitive Science* 11: 259–281.
- Dascal, M. 1989. "On the roles of context and literal meaning in understanding". *Cognitive Science* 13: 253–257.
- Dascal, M. 2002. "Language as a cognitive technology". *International Journal of Cognition and Technology* 1(1): 35–89.
- Dennett, D. C. 1991. *Consciousness Explained*. Boston: Little, Brown.
- Dummett, M. 1989. "Language and communication". In A. George (ed), *Reflections on Chomsky*. Oxford: Blackwell. Reprinted in M. Dummett, *The Seas of Language*, Oxford: Clarendon Press, 1993.
- Fodor, J. A. 1998. *Concepts*. New York: Oxford University Press.
- Fodor, J. A. 2001. "Language, thought and compositionality". *Mind and Language* 16: 1–15.
- García-Carpintero, M. 2001. "Gricean rational reconstructions and the semantics/pragmatics distinction". *Synthese* 128: 93–131.
- Gauker, C. 1994. *Thinking Out Loud*. Princeton, NJ: Princeton University Press.
- Harman, G. 1975. "Language, thought, and communication". In G. Harman, *Reasoning, Meaning, and Mind*. Oxford: Clarendon Press [1999], 166–182.
- Jackendoff, R. 1987. *Consciousness and the Computational Mind*. Cambridge, MA: The MIT Press.
- Jackendoff, R. 1997. *The Architecture of the Language Faculty*. Cambridge, MA: The MIT Press.
- Katz, J. J. 1977. *Propositional Structure and Illocutionary Force*. New York: Crowell.
- Levinson, S. 1997. "From outer to inner space: Linguistic categories and non-linguistic thinking". In E. Pederson and J. Nuyts (eds), *Language and Conceptualization*. Cambridge: Cambridge University Press, 13–45.
- Lucy, J. 1992. *Language Diversity and Thought*. Cambridge: Cambridge University Press.
- Mithen, S. 1995. *The Prehistory of Mind*. London: Thames and Hudson.
- Recanati, F. 1995. "The alleged priority of literal interpretation". *Cognitive Science* 19: 207–232.
- Recanati, F. 1997. "Can we believe what we do not understand?". *Mind and Language* 12: 84–100.
- Recanati, F. 2001. "What is said". *Synthese* 128: 75–91.
- Recanati, F. 2002a. "Does linguistic communication rest on inference?". *Mind and Language* 17: 105–127.
- Recanati, F. 2002b. "Literal/nonliteral". *Midwest Studies in Philosophy* 25: 264–274.

- Recanati, F. 2003. *Literal Meaning*. Cambridge: Cambridge University Press.
- Sellars, W. 1969 "Language as thought and communication". *Philosophy and Phenomenological Research* 29: 506–527.
- Spelke, E. and Tsivkin, S. 2001a. "Initial knowledge and conceptual change: space and number". In M. Bowerman and S. C. Levinson (eds), *Language Acquisition and Conceptual Development*. Cambridge: Cambridge University Press, 70–97.
- Spelke, E. and Tsivkin, S. 2001b. "Language and number: a bilingual training study". *Cognition* 78: 45–88.
- Sperber, D. and Wilson, D. 1986. *Relevance: Communication and Cognition* (2nd edition, 1995). Oxford: Blackwell.
- Vicente, A. and Martínez-Manrique, F. In press. "Semantic underdetermination and the cognitive uses of language". *Mind and Language*.
- Waismann, F. 1951. "Verifiability". In A. Flew (ed). *Logic and Language*, 1st series. Oxford: Blackwell, 17–44.

Authors' addresses

Fernando Martínez-Manrique
 Dpto. de Filosofía
 Fac. de Psicología
 Univ. de Granada
 18011 Granada,
 Spain
 Email: fernan_martinez@yahoo.com

Agustín Vicente
 Dpto. de Filosofía
 Fac. de Filosofía y Letras
 Univ. de Valladolid
 47011 Valladolid, Spain
 Email: gusa@euskalnet.net

About the authors

Agustín Vicente is Assistant Professor of philosophy at the University of Valladolid. His work has been mainly centered on the philosophy of mind, especially on its metaphysical side. He has published in journals such as *Erkenntnis*, *Minds and Machines*, and *Crítica*. He is currently working on the relationship between language and thought and on the problem of mental causation.

Fernando Martínez-Manrique is "Ramón y Cajal" Researcher at the University of Granada. His work is focused on foundational problems of cognitive science, particularly the symbolic-connectionist debate. He has published in *Minds and Machines* and *Theoria*. Currently he works on the relations between language and thought and its relativistic implications, as well as on inter-theory relations in cognitive science.

