Semantic Underdetermination and the Cognitive Uses of Language

AGUSTÍN VICENTE AND FERNANDO MARTÍNEZ-MANRIQUE

Abstract: According to the thesis of semantic underdetermination, most sentences of a natural language lack a definite semantic interpretation. This thesis supports an argument against the use of natural language as an instrument of thought, based on the premise that cognition requires a semantically precise and compositional instrument. In this paper we examine several ways to construe this argument, as well as possible ways out for the cognitive view of natural language in the introspectivist version defended by Carruthers. Finally, we sketch a view of the role of language in thought as a specialized tool, showing how it avoids the consequences of semantic underdetermination.

1. Introduction: Semantic Underdetermination

A number of authors have argued forcefully for a thesis that is generally known as 'semantic underdetermination' (see Bach, 1994; Sperber and Wilson, 1986/95; Recanati, 2002). According to this thesis, there are many expressions of a natural language for which semantic interpretation cannot determine by itself what is said by a sentence containing such an expression. Rather, in order to ascribe a definite meaning to a sentence, it is necessary (in many cases) to take contextual factors into account. There are some typical examples that exhibit this phenomenon. Thus, what is said in sentences such as 'John's car is empty', 'the chair is made of oak wood', 'that dog looks dangerous', or 'all the students have gone on strike' is semantically underdetermined. 'John's car' can mean a number of things: we may be talking about the car John owns, about the car John designed, about the car John drives, etc. What particular chair we are referring to in the second example is not specified until we know something about the context in which the sentence was uttered. Demonstratives like 'that', Recanati remarks, lack a 'meaning rule' such as the one governing the interpretation of 'I', and acquire a definite meaning only when pragmatic factors are taken into account. Last, in order to know the

Address for correspondence: Departamento de Filosofía, Facultad de Filosofía y Letras, Universidad de Valladolid, 47011 Valladolid, Spain; Departamento de Filosofía, Facultad de Psicología, Universidad de Granada, 18011 Granada, Spain.
Email: gusa@euskalnet.net; fernan_martinez@yahoo.com

truth conditions of 'all the students have gone on strike' it seems that we need further information about the domain of the universal quantifier.

The phenomenon these examples illustrate also involves the failure of the principle of compositionality in natural languages. They seem to show that the meaning of a whole (a sentence) is not determined by the meaning of its parts because some of its parts lack a definite meaning. That is, the principle breaks down because there are terms in natural languages that suffer intrinsically from semantic underdetermination, such as comparatives, definite descriptions, demonstratives or universal quantifiers. We cannot ascribe a semantic value to terms of these types without pragmatic clues. We cannot construe the meaning of all sentences by semantic means alone. Perhaps one could say that this does not show that the principle of compositionality fails, for once we have decided what meaning the 'problematic' terms have, the meaning of the whole does depend solely on the meaning of its parts. However, the decision about what these problematic terms mean is not made independently, outside the context of the whole utterance.

The semantic underdetermination (henceforth SU) of natural language (henceforth NL) sentences has consequences that go beyond its usual application to discriminate between pragmatic theories. In this paper we want to examine one of those consequences: the question whether we do or do not think in NL. The structure of the paper is the following: First, we will present the main lines of the cognitive view of language, according to which it is possible to think in one's NL, paying special attention to Peter Carruthers's introspectivist proposal. Then, drawing upon an argument by Fodor, we will show how the link between SU and non-compositionality undermines Carruthers's project. However, we will identify a strong and a weak argument from SU to non-compositionality, and we will argue that Fodor's version is too weak to make the point. Having identified in Recanati a stronger argument, the cognitive view of NL seems to have a strategy to avoid its conclusion, provided that two conditions are met: that there is some layer of meaning in which a NL sentence may appear as compositional, and that this layer of meaning is psychologically realistic. The first condition might be met by exploiting the possibilities of explicitation that NL itself offers. However, even if NL is explicitable in the required sense—something that is disputable—we will contend that this property runs counter to the introspectivist evidence that gives psychological plausibility to Carruthers's proposal. In addition, a final way out in which psychological reality is sustained by unconscious NL sentences, would reverse the direction of the evidence on which he intends to ground his theory. Finally, we will sketch a (milder) view of the role of language in thought that, while allowing a limited use of NL in cognition, as a specialized tool, does not have to face the consequences of SU.

2. The Cognitive View of Language

After at least three decades of anti-(linguistic) relativism and a reversal of the linguistic turn, some authors are beginning to defend a cognitive use of natural

language. For instance, Dennett (1991) views language as a virtual architecture on which conceptual thought runs; Bickerton (1995) argues that only through the evolution of language was conscious, abstract thought possible; Elizabeth Spelke (see Spelke and Tsivkin, 2001a, 2001b) has proposed that our mother tongue may be acting as an intermodular *lingua franca*; Stephen Mithen (1995) collects archeological evidence to advance a similar proposal, while Andy Clark (1998) suggests that we use natural language as a way to approach our thoughts in a reflexive way, to 'contemplate' our own thoughts, so to speak. But it is Peter Carruthers (1996, 2002) who has defended the strongest position about this topic.

According to Carruthers, we make use of natural language as a *vehicle* of some of our thoughts. That is, some thoughts that we have are, literally, the content of sentences of a NL. Carruthers has provided two kinds of arguments for this position. One of them (Carruthers, 1996), which will occupy us here, is that our own introspection reveals that we do use language when we think. Very schematically, his position is that (a) data from introspection reveal that sometimes we think in a NL, and (b) that we must begin by taking these data from introspection for what they seem. Moreover, he argues that we should go beyond what introspection reveals—that is, that we codify linguistically episodic conscious thoughts—and admit that we also use NL to codify latent thoughts and unconscious token-thoughts that belong to the same types of which we know conscious token-thoughts. His second argument (Carruthers, 2002) draws on experimental evidence by Spelke and others to conclude that natural language is the vehicle of non-modular, non-domain specific, conceptual thinking which integrates the results of modular thinking. Even though language constitutes a module itself, its being both an input and output system would grant it a privileged position to combine different kinds of information. These sorts of considerations are beyond the scope of this paper. Nevertheless, we think that the argument from semantic underdetermination that we are about to present is general enough to affect Carruthers's latest position, regardless of his new modularist backbone, because it is based on a condition that thought must, and language cannot, meet. Furthermore, we will argue later (see section 7) that language might still be an intermodular integrator, even though it does not work as a vehicle of thought.

There are several types of counterevidences and counterarguments that might be alleged against Carruthers's thesis. For instance, he maintains that the use of language in cognition is necessary, at least in order to perform some definite cognitive functions. However, even though linguistic deficits are often associated with cognitive deficits in aphasic disorders, it seems that there are aphasics who have experienced no noticeable deterioration in any cognitive function (see Varley, 1998). There are also illustrious arguments he has to face, like Fodor's (1975) language acquisition argument for the language of thought, according to which there is no possible way to learn a language except by translating it into an already existent language. But the strongest argument against a position like Carruthers's comes from the apparent semantic underdetermination of natural languages. Cognition, this argument says, requires a semantically precise and compositional

instrument or vehicle. As NLs are largely semantically underdetermined, they cannot be a vehicle of cognition.

There are a number of authors who have defended versions of this argument. Here we will single out Steven Pinker (1994) and Jerry Fodor (2001), who have advanced two completely different versions of the argument. Good expositions of one version or the other, however, can be found elsewhere (e.g. McDermott, 1981; Levinson, 1997).

3. Turing Machines

Drawing on ideas advanced by Drew McDermott (1981), Steven Pinker, in *The Language Instinct* dedicates one chapter to arguing against the cognitive use of language. After presenting some evidence that seems to prove that language is not necessary for cognition, he discusses why a natural language could never be an instrument of thought. His thesis is that the human mind/brain works like a Turing machine, and that no NL can function as a language for the use of a Turing machine. As Fodor repeatedly remarks, Turing's great discovery was to provide a means by which rational transitions between thoughts could be explained and implemented. T-machines attend solely to syntactic, or formal, properties of a language but are capable of preserving truth. However, in order to do that, they need a special kind of language. They need a highly regimented, unambiguous and compositional language. They cannot work with a language that does not have a distinct term for every distinct concept, or whose terms and sentences express different concepts and propositions on different occasions of occurrence. A T-machine will conclude from 'all students went on strike', and 'all strikers try to destabilize the Government' that all students try to destabilize the Government. From here, and the information that Peter is a student, it will conclude that Peter tries to destabilize the Government. However, Peter may well be studying at a school that has never been on strike (i.e. Peter may well be outside the domain of the quantifier in 'all students went on strike').

Let us say that a T-machine requires a language that is informationally exact. By this we mean that its sentences express or codify all the information (no more, no less) that they are intended to express. It is a language that codifies every piece of information necessary to reach the proposition that the sentence is intended to express. In the absence of a language like that, a T-machine will be unable to work: it will not preserve truth, deriving false conclusions from true premises. Thus, a language that suffers from semantic underdetermination, so that it requires contextual information to get the proposition a sentence is intended to express, cannot afford the code that a T-machine demands. Given that, arguably, natural language sentences are subject to SU, then NLs are hopeless instruments of thought.

One may argue (in fact, both McDermott and Pinker do) that this shortcoming of natural languages is not an accidental feature. It derives from their primary or proper function, i.e. being instruments of communication. It is this fact that

explains that it cannot be a derived function of language to be the means by which we think. When we 'speak our thoughts' we have limited time, in which not only do we have to say what we want to say, but we also have to attract and keep the attention of our audience. So it is not surprising that we make as much use as we can of contextual factors in order to communicate our thoughts. We exploit the context, and omit all that information that can be easily gathered from it. There is no reason, then, why natural languages should be informationally exact in the sense explained above. Rather, the opposite seems right: natural languages should most probably be inexact.

4. Thoughts and Contents

Pinker's argument may be resisted because of its commitment to Turing machines. It is an argument that will not convince the enemies of the computational theory of mind, or those that conceive of computational minds as very different from T-machines.[1] However, there is a more general argument from semantic underdetermination that seems to establish the same point Pinker's argument was designed to make. The last author who has presented a variant of this other argument, quite surprisingly, is Jerry Fodor (2001). We will start with the argument, and then explain why it is surprising that Fodor has endorsed it.

Let us assume again that the semantics of natural language is underdetermined and, hence, noncompositional. Now the question is: can a language of thought (or more generally, a vehicle of thought) have an underdetermined noncompositional semantics? According to Fodor, the answer is 'no', and the reason is that a thought cannot be inexplicit (or inexact) about its content because, he says, 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 LOT sentence can have *only one content*.

Let us explain this with a little more detail. Fodor (1998a) conceives of thoughts as being composed of conceptual atoms. These ultimate components of Mentalese, unlike NL words, are context-independent. 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 a unique content. But this is generalizable to whole thoughts: a thought has also a unique content that is determined by the contents of its component concepts. Thus 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 express exactly[2] the content it is meant to express. 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, it cannot be a vehicle of thought.

We say that it is surprising that Fodor endorses this argument because it is really difficult to match it with his previous views about the topic. It is true that, for many years, he has argued that there must be a language of thought precisely because thought is compositional. His most popular argument has been that as thought is systematic and productive, it must have a compositional semantics. In some places (Fodor and Pylyshyn, 1988), he has just assumed that thought is systematic and productive, but there is at least one place where he has tried to justify it (Fodor, 1987, Appendix). The complete argument for the language of thought presented there can be summed up in this way:

    (i)   natural languages express thought;

    (ii)   natural languages are productive and systematic;

    (iii) (from (i) and (ii)) the instrument of thought must be just as productive and systematic as natural languages are;

    (iv) the best explanation for the properties of productivity and systematicity is a compositional semantics;

    (v)   (from (iii) and (iv)) the instrument of thought has a compositional semantics.

An intriguing consequence of Fodor's endorsement of the view that NLs are semantically underdetermined is that he has to drop premise (iv) above. In effect, NL provides a case where systematicity and productivity are not explained by the existence of a compositional semantics (because, despite semantic underdetermination, NLs are systematic and productive). Hence, one may wonder why it is that the instrument of thought has a compositional semantics. In his (2001), Fodor starts by stating that, after all these years, it has been shown that thought must be compositional. However, if that had been shown, it would have been by means of an argument that semantic under-determination puts in jeopardy.

---

[2] We want to avoid a confusion with regard to the notion of 'exact thought'. This does not mean that there cannot be a certain element of vagueness in thoughts, for instance, because its truth conditions are indeterminate. The possibility of inherently vague concepts (say, BALD) is something that we take to be an empirical question. But the existence of a vague concept does not imply that the thought where it occurs needs further interpretation (say, to determine exactly how many hairs we are referring to): the thought may be exact in the sense of expressing a unique complete content (fixed by the component elements plus rules of composition), even though that content may be vague.

So what Fodor says in his (2001) runs counter to what he has been saying for a long time, although it is a point that it is not clear that he has noticed. In any case, despite this element of surprise, it is possible to distinguish in Fodor's (2001) an independent argument for the compositionality of the language of thought. Thought must be compositional because a thought is explicit about its content. We take this to be a good argument for the informationally exact character of the vehicle of our thoughts, and its consequent compositionality. In fact, more than an argument, it is what we could call a 'phenomenological realization': we know that thoughts are complete, that we do not have to go through a process of interpretation to get the 'real' information that one of our thoughts is trying to provide. If this is indeed a good proof that the instrument of thought is exact, then it is thereby a good proof that natural language cannot be an instrument of thought (since, we have assumed, natural language is semantically underdetermined).

5. Two Arguments from SU to Non-Compositionality

So far, we have shown how the semantic underdetermination of NL undermines Carruthers's project, because it entails the non-compositionality of NL, and without compositionality it cannot be a vehicle of thought, as Carruthers contends. However, we want to refine this discussion and distinguish two arguments from SU to the non-compositionality of NL. They are only slightly different, so it is easy to mistake one for the other. Yet, we take it that only one of them provides a strong argument.

The strong argument 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 Underdetermination*) 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) can be generalized to most NL sentences.

Hence:

(d) Compositionality does not hold for most NL sentences.

The weak argument differs in premise (b), and (c) is subsequently modified:

(b')   There are cases in which it is not possible to obtain the speaker's meaning expressed by sentence S solely from the semantic value of S: it is necessary to add pragmatic factors.

(c')   Point (b') can be generalized to most NL sentences

Hence (d).

Why is this a weak argument against the compositionality of NL? Because NL sentences can still be compositional *with respect to a certain meaning*. Call this the semantic or literal meaning. What the weak argument says is that literal meaning does not determine speaker's meaning: the value of the latter cannot be obtained from the semantic value of the parts of the former (plus rules of composition). But it does not rule out that literal meaning *itself* can be obtained compositionally. So if the defender of the cognitive view could show that there is room in the mind for this sort of literal meaning, he might have a chance to support the thesis that NL is a compositional instrument of thought.

In contrast, the strong argument denies that even the literal meaning can be compositional. What it says is that most NL sentences *do not have* a literal meaning. Since neither lexical components nor rules of composition have definite values *prior* to the addition of pragmatic, contextual factors, the conditions that make compositionality possible are destroyed root and branch.

In our view, Fodor's argument belongs to the weak type. The kernel of his argument is expressed in the following lines:

> But, as a simple matter of fact, in the general case, sentences are remarkably *in*explicit with respect to how the thoughts they express are put together. So either the content of the thought is different from the content of the sentence that expresses it, or the sentence isn't compositional. I take it that the first disjunct is preposterous; so I take it that the second disjunct must be true. (Fodor, 2001, p. 12).

Fodor is assuming that there is a single layer of meaning, the speaker's meaning. If this were the case, then the transition from premise (b') to (d) would be warranted: if the speaker's meaning is not semantically attainable, and there is no other meaning to consider, then meaning can never be obtained compositionally. However, Fodor does not explain why the difference in content that he mentions is 'preposterous'. On the contrary, it is a typical pragmatic move to distinguish the content of a sentence from the content of the speaker's thought it is used to express. We can safely assume, with Recanati (2001), that there are at least three levels of meaning involved in an utterance: sentence meaning, what is said, and what is implicated. Nothing that Fodor says rules out the possibility that, even though the latter two meanings cannot be obtained compositionally from the sentence, sentence meaning itself can be compositional. If this meaning had psychological reality, then Carruthers's position would gain a foothold.

In contrast, Recanati's view contains an argument that is an instance of the strong type.[3] Let us review briefly what goes on, according to Recanati (2001), when you hear a sentence in your natural language. To get 'what is said' by a given sentence involves, at one and the same step, decoding the meaning of the sentence, filling in the referents of the indexical expressions and going through the process of 'free enrichment';[4] after that, you draw all the implicatures needed to reach the intended meaning of the expression. So, as a hearer, you go through the process of interpreting somebody's utterance in order to arrive at a content, at a thought, that corresponds, at least, to what is said, and, possibly, to what is implicated.

Thought is that place where everything becomes clear, so to speak. It is in that essential aspect that the realm of thought is different from the realm of linguistic meaning. We do not have indexical thoughts that we have to interpret, and there is no information coming from the context that you have to add to the thought in order to make it complete. The contrast between thoughts and propositions of a natural language goes like this: thoughts are psychologically real entities where all the information comes together—the information conveyed by a statement plus that contained in the context; propositions (i.e. the purely semantic or literal meaning of an expression) are theoretical abstractions that are never entertained in the process of interpreting a given utterance. Hence, the real content of a sentence, what is obtained when it is uttered, is always the product of pragmatic interpretation that takes place in the interpreter's head. Natural language cannot by itself provide such a content.

Recanati's examples of demonstratives like 'that' in 'that dog', or syntactic relations like ('s) in 'John's car', are intended 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, then most

---

[3]   A referee has shown concern for the relevance of the distinction between a strong and weak argument, pointing out that both Recanati and Fodor claim that SU is a structural phenomenon of NL. Hence in both views NL would be precluded from being the language in which we think. We agree that this is the aim of both arguments, yet it is our contention that Fodor's argument can only work on the assumption that the speaker's meaning (i.e. the thought the interpreter gets) and the sentence meaning are one and the same (therefore his dictum (Fodor, 2001, p. 13) that 'English hasn't got a semantics; the study of its semantics is the study of the semantics of thought'). This assumption is heavily contested by pragmaticians who are at pains to tell the different layers of meaning apart. Recanati's argument is stronger because it reaches its conclusion taking into account these pragmatic distinctions. Our claim, and this will be the whole point of section 6 below, is that the strategy open to the cognitive view of language is to try to exploit the existence of those layers, so as to argue that one of them is psychologically realistic and corresponds to a compositional NL. We will show, in three stages, that this attempt is unsuccessful.

[4]   Free enrichment is the process by which we add extra elements to an expression, according to a particular context. For instance, it is by free enrichment that one interprets that 'I've had breakfast', said by someone as a response to 'are you hungry?', means that she has had breakfast that very morning (and not that she has had breakfast at least once in the past).

of them do not meet the conditions for being compositional, even at the level of sentence meaning. If Recanati's examples are correct, then SU denies the compositionality of NL *tout court*, not at a particular level. This is the reason, then, why NL sentences cannot be a vehicle of thought: their semantics cannot furnish complete mental contents. Mental content is obtained *from* NL sentences but it cannot be *equated with* those sentences.

6. Explicable NL Sentences and Introspective Data

As we hinted in the previous section, there seems to be a strategy for Carruthers's cognitive view of language to escape the argument from semantic underdetermination. This strategy is based on two conditions. First, he would need to show that there is some layer of meaning in which a NL sentence may appear as compositional. Second, he would have to show that this layer of meaning is psychologically realistic. In this section we want to argue that even if it were possible to construe NL so as to meet the first condition, by doing so Carruthers would undermine the introspectivist evidence on which he establishes the second one. We will do this in three stages. First, we will examine a possible construal in which NL sentences may be *explicable*: a sentence S is explicable when there is another NL sentence S★ such that (a) S★ has the same content as S, and (b) S★ is explicit about its content, in the sense of being semantically determinate. We take that even though there are grounds for being dubious about this construal, there is still no knock-down argument against it. We will show then that this construal runs counter to the introspective evidence that, according to Carruthers, motivates the case for the psychological reality of NL as a vehicle of thought. Finally, we will reject a last-ditch effort to sustain this psychological reality in terms of unconscious structures, arguing that it reverses the direction of the argument with which Carruthers began.

6.1 Explicable Natural Language Sentences

We began this paper by giving some examples, which show that sentences of a natural language are semantically underdetermined. These sentences, it can be conceded, are in fact paradigmatic, since most sentences that are uttered are similarly in need of information coming from the context in order to get their truth-conditional meaning (premise (c)). However, one may wonder whether this is not a redeemable feature of these sentences. Admittedly, most sentences that are uttered are not complete or exact, but are they not completable by linguistic means alone? To see what we mean by this, let us begin by distinguishing three possible kinds of sentences, subject to different kinds of underdetermination.

First, we have sentences where there seems to be a real semantic underdetermination, such as 'John's car is empty' or 'that dog looks dangerous'. It can be argued that (apart from the problematic terms 'empty' and 'dangerous') the meaning of these sentences is underdetermined until we solve what relation there is between

John and his car in the first case, and what 'that' stands for in the second. Since these sentences do not have a definite meaning until pragmatic factors are taken into account, they do not even have a minimal, literal meaning to begin with (Recanati, 1995, 2001)

Second, we have cases such as 'the chair is made of oak wood' or 'all the students have gone on strike'. Here, it may be said, there are definite meanings: the problem is that such meanings are not the meanings those sentences are conventionally used to convey. According to the Russellian analysis of definite descriptions, 'the chair is made of oak wood' means that there exists a single chair such that it is made of oak wood. But at most, what we want to say with that sentence is, say, that there exists a single chair *in this room* such that it is made of oak wood.[5] That is, in these instances we have compositionality, yet the composed meaning is an undesired meaning, a proposition that is patently false and unintended. These sentences have a literal meaning, but it is different from the conventional meaning that the speaker intends to convey.

Third, there might be sentences with definite meanings that correspond to the meaning that the sentence is conventionally used to convey. 'I am Peter' may be one example. These sentences would not be subject to semantic underdetermination and they appear to be compositional. Their literal and conventional meanings are one and the same.

Now, given that an instrument of thought must be compositional, only type-2 or type-3 sentences are fit for that job. Of these two kinds, only type-3 sentences would be capable of expressing in a explicit, compositional way the speaker's meaning. Type-2 sentences can be the vehicle of their compositional meaning, but not of the speaker's meaning that they conventionally express. In other words, if we can have a thought in a type-2 natural language sentence, then (inasmuch as we grasp the meaning that the speaker conventionally conveys) we must have 'in another place' a thought that is *not* a type-2 NL sentence and that expresses (in a complete, exact way) the speaker's meaning. This latter thought *might* be a type-3 sentence.

When we wonder whether NL sentences can be completable by linguistic means alone, we are in fact asking whether there is a translation from a type-1 or a type-2 sentence into a type-3 one. For instance, it is true that we say 'John's car is empty', but we could just as well say 'the car owned by John is empty in the sense

---

5   There are some attempts to analyze those sentences containing (overt or covert) quantifiers (type-2 sentences, say) as sentences that have a hidden indexical component that would point to the adequate domain restriction (see specially Stanley and Szabo, 2000). And there might be someone disposed to treat these same examples as instances where Gricean implicatures would lead you to grasp the intended meaning via the literal meaning that delivers the patent falsity. Yet, this is not very likely to succeed. As Recanati points out, Gricean implicatures are derived from available propositions, whereas when listening to 'the chair is made of oak wood', we do not consciously entertain the proposition that there is a single chair such that it is made of oak wood, and then go on to reflect on whether the speaker is intending to communicate that or not.

that there is no person in it'. One might claim that the latter sentence suffers, at most, from type-2 underdetermination. There are still problematic terms, such as 'the', but the problem now would not be that the sentence does not have a definite meaning. Rather, it would be (at most) that its literal meaning does not coincide with its conventionally intended meaning. Going a step further, maybe type-2 sentences are translatable into type-3 sentences, that is, sentences that literally express the conventional meaning (that would correspond to Recanati's 'what is said'). The resulting sentence would possibly be very long and utterly unmanageable, but it would be a NL sentence nevertheless. The reasons why one does not use this precise translation are broadly Gricean: a speaker should not use more words than she needs.

To put it in a different way, as we said in section 3, McDermott (1981) and Pinker (1994) have it that the inexact character of natural languages derives from their being instruments of communication. The limited time that we have when we speak, and the many things we have to do in that limited time, explains that when we utter a sentence we exploit the context as much as we can. However, this is a good explanation for why most of the sentences that we utter *are* inexact. It is not a good explanation for why natural languages *should* be inexact. That is, McDermott's and Pinker's thesis explains why we choose to say 'John's car is empty' instead of saying 'the car owned by John is empty in the sense that there is no person in it', or something even longer. But, of course, it does not explain why a natural language does not have exact sentences. It could not explain it, because it seems that NLs are capable of producing exact sentences; moreover, that they are capable of completing the inexact sentences that we utter.

Now, the strong argument from SU to non-compositionality demands type-1 (or genuine) underdetermination: this is the kind of SU that premise (b) states (in contrast, premise (b') in Fodor's weaker argument demands only type-2 underdetermination). But if this SU is not an irredeemable feature of natural language, then NLs *could* be the compositional instruments of thought. The translations to type-3, complete, exact sentences are never uttered in conversation, but they could be carried out 'in the head'

However, it is unclear whether the proposed translation to type-3 sentences is possible, *even in principle*, because it is unclear whether type-3 sentences can exist at all. Following Recanati (2001, 2002) it can be argued that semantic underdetermination is not merely a result of the conditions in which communication occurs: it is related to the nature of linguistic construction itself. He suggests, for instance, that there is constructional underdetermination in Adjective + Noun constructions like 'red pen'. These constructions would be dependent on wide contextual factors in order to have a definite meaning. In addition, following Waismann (1951), he suggests that sentences with empirical predicates are subject to SU. Those two cases should be enough to dispute the claim that SU is a redeemable feature of language, at least for vast portions of language: it is unlikely that we can explicitate an Adj + Noun construction in a structure that is not dependent on context, and it is, of course, impossible to translate an empirical predicate into a

non-empirical one. Moreover, there are possibly further cases of underdetermined structural elements of language, such as prepositional constructions, adverbial phrases, and many more. Altogether we can regard the existence of type-3 sentences as questionable: the interpretation of any sentence seems to demand a context.[6]

Even though we are sympathetic to this general line, we think that Recanati's observations fall short of proving that strong conclusion. To this end, it would be necessary to provide an in-depth analysis of the expressive possibilities of languages: perhaps there are, within the boundaries of a NL, resources to construct the required type-3 translations, even if the resulting constructions are too alien to be normally uttered sentences.[7] Still, the defender of the cognitive view of language must show that these constructions can indeed constitute a vehicle of thought, and we can draw further considerations from Recanati (1995, 2001) to cast a shadow over this project.

Recanati argues against the existence of a stage of literal interpretation of an utterance. In the standard view, the literal interpretation of an utterance is computed first, and any nonliteral interpretation is obtained inferentially from it. In his alternative model, the literal interpretation of constituents is accessed before any other interpretation, but not so the literal interpretation of the whole. It is possible for a nonliteral interpretation to be reached first, provided that its elements are

---

[6]   Even our attempt 'I am Peter' is a dubious example of a type-3 sentence. There are pragmatic contexts that yield different interpretations, e.g., suppose that Peter is under a death sentence and I am expressing my support by uttering that sentence. Even if there is a conventional meaning for an expression, we need to know that it is the speaker's intention to utter that expression in the conventional sense, and this knowledge is dependent on the wide context.

[7]   A referee has suggested that we have read Carruthers's claims in an unnecessarily strong way. The claim that NL sentences are the medium of (some) thought is compatible with the idea that NL sentences, together with other cognitive elements, carry the content of the thought. That is, it does not imply that NL sentences have to exhaust the content of the thought. Such extra cognitive elements can be perceptual representations, for instance, when one entertains in inner speech the sentence 'that [perceived lamp-stand] is about to fall over' (referee's own example). One may think that if it were not for this strong reading, we would not get to the conclusions we are trying to establish at this point. For it is one thing to say that NL cannot provide adequate translations to type-3 sentences and another that NL with the aid of images cannot do. Perhaps NL alone cannot remove the underdetermination of 'that lamp-stand is about to fall over', but NL-plus-images, as the example shows, can do it. Now, we take it that we do not need this strong reading of Carruthers. NL-plus-images can be an explicit medium in instances like 'that [perceived lamp-stand] is about to fall over, and so remove the indeterminacy introduced by the demonstrative in 'that lamp-stand is about to fall over'. However, it is doubtful that such a mixed vehicle could remove all those indeterminacies that do not involve demonstratives, such as the one introduced by the genitive in 'John's car is empty' or the one introduced by 'here' in 'it is raining here'. As for the next point to be discussed, we also doubt that we have introspective access to explicit sentences made of NL and images. That is, the point we want to make in the next section stands both for NL-only sentences and NL-plus-images sentences. We do not usually introspect entire sentences, and when we do, they are not free of ambiguities or underdeterminations.

associatively derived, by spread of activation, from the constituents that were initially accessed. In addition, contexts can interact in such a way that apparently optional processes, like nonliteral interpretation, have to operate before the mandatory processes that presumably would recover the literal interpretation.

To round out the view, we can add Recanati's (2001) arguments against the psychological reality of minimal propositions (the most plausible account of literal interpretations). If Recanati is right, the notion of a literal interpretation of a NL sentence is a theoretical artifact, not a psychologically real entity. But if literal interpretations cannot be entertained in mind, then type-3 sentences are not a realistic candidate for instrument of thought. At most, they would be linguist's fictions: sentences that one can construct only by resort to linguistic theory, not sentences that the cognitive system can actually compose.

Yet, there is a weaker reading of Recanati's arguments. First, they count against the reality of literal propositions *only* with respect to semantically underdetermined sentences. So if it is indeed the case that language is explicitable, his considerations do not rule out that an interpreter can extract a literal interpretation from explicited sentences, which are not semantically underdetermined. Second, Recanati's view rejects the existence of an early stage of decoding whose result is a literal or minimal proposition. The psychological unreality of this notion has to do with its unavailability to consciousness. Putting it in our terms, what it discards is the idea that, in obtaining the speaker's meaning, the interpreter goes through a process of extracting an intermediate type-3 sentence (that corresponds to the literal interpretation of the uttered sentence). But it does not discard the idea that the speaker's meaning itself can be entertained *at some final point* in the form of a type-3 sentence (that expresses in a complete, exact way the speaker's meaning).

Going back to Carruthers, we have tried to provide him with a way out of the SU argument: if NL sentences can be rendered into literal, compositional translations, then they can constitute a vehicle of thought. We have drawn upon Recanati's considerations to challenge this possibility. On the one hand, it is possible that those translations do not exist. On the other, even if they exist, it is possible that there is no room for them in the mental economy. We have concluded that the jury is still out with regard to both possibilities. So, in the absence of stronger arguments, natural language might be explicitable and play the cognitive role that Carruthers wants it to play. However, we want to show next that the explicitable character of NL sentences is not the kind of property on which Carruthers can base his position, because the evidence with which he supports it runs counter to the presence of explicit, exact sentences.

6.2. What Kinds of Sentences Does Introspection Reveal?

Let us recall the argument Carruthers presents in favor of the cognitive use of language as a vehicle of thought. To put it in a nutshell, what he claims is that we should take the data from introspection at face value, and that introspection reveals many uses of language in thought (or, rather, as thought). When we write or read,

we 'tell ourselves' the words that we write or read. Also, when we write it can be said that we think writing: the process of writing and the process of thinking what is to be written take place at the same time, and contents simply flow from what is written to what is thought and vice versa. We also have the experience of thinking in a foreign language that we have been learning, or of going through a dialogue with an imaginary interlocutor. Finally, when we perform an action we are not good at, we sometimes discover ourselves rehearsing the verbal instructions we have been told before.[8]

If natural language were an instrument of thought, then (as discussed in section 4) the sentences we entertain should be informationally exact. That is, by means of introspection we should discover 'in our heads' token sentences of the sort 'the car owned by John is empty in the sense that there is no person in it', rather than underdetermined token sentences like 'John's car is empty'. However, except in the case of reading and writing, we do not usually entertain complete sentences. Moreover, it seems that the linguistic items that we experience are specially fragmentary, i.e. even more underdetermined than the sentential or quasi-senten-tial items of spoken language. We think this is something a simple reflection on our phenomenology (or on literature) will reveal: when we feel that we are thinking in language most times we do not discover entire sentences but two or three words that sum up a complete proposition. In fact, often we have the feeling that we are 'speaking to ourselves', that is, that we are in a situation where linguistic economy can be maximized because speaker and audience have a maximal mutual knowl-edge (why one speaks at all in such a case is something that we will try to explain below). Internal monologues in the novel of the nineteenth century were com-posed of entire sentences, but it is the novel of the twentieth century, with its fragmentary internal monologues, which is more realistic in this sense.

This point is reinforced by some studies carried out by W. Frawley (1997). As has been said, one of the occasions when we realize that we are making an internal use of language is when we rehearse instructions. This seems to be a case analogous to that described and studied by Vygotsky (1934/1962) of very young children going through monologues. Now, Frawley's research shows that the language used in those cases is of a special kind. For instance, it contains more verbs in infinitive than tensed forms, and it employs a special kind of predication, that he calls psychological predication, which keeps most referential elements, while removing many functors. Frawley labels this private speech as 'language *for* thought', and regards it as the bridge between the external context and the internal computa-tional architecture (i.e., the language *of* thought). If his analysis is correct, the language for thought should not be equated with the internal code, or with public language: it is a specialized subsystem involved in the control of thought and

---

[8]   Carruthers's idea could be attacked at the root, denying that introspection is a reliable source of evidence. We will leave this criticism aside, and will concentrate on how sustainable his position is, supposing that introspection is reliable. What we contend is that the data do not show what Carruthers claims.

behavior, and in the deployment of metaconsciousness. The origin of this sub-system may lie in the acquisition of natural language. Vygotsky hypothesized that young children, after some time speaking aloud, internalize that use of language, and go through the same kind of monologue, only this time in an internal way. If this is correct, then it is probable that when we adults give instructions to ourselves quietly, so to speak, we do so in the peculiar kind of language for thought Frawley proposes. Even if this language derives from NL, it is distinct from NL itself. Hence, it is not surprising that the sentences of this language are even less compositional and less explicit about their content.

To sum up, introspection does not reveal to us the kinds of NL sentences that are capable of expressing complete, exact thoughts, even if these sentences exist.[9] Introspection reveals fragmentary, patently non-compositional sentences, that cannot fulfill the requirements of Carruthers's view. Still, there is another way to salvage this view: maybe there is another place, within the boundaries of mind, where we can look for those explicit NL sentences.

6.3. Unconscious Natural Language Sentences

As we said in section 2, Carruthers (1996) extends natural language toward the realm of latent and unconscious thoughts. Whenever we have grounds to maintain that an unconscious thought is of the same kind as conscious thought, we can infer that the vehicles of both are the same. Hence, if conscious thoughts are framed in NL, we can conclude that unconscious thoughts will share this code. Carruthers (2002) puts this point in a different way. He claims that the linguistic representations used for intermodular integration are logical form (LF) representations. Whenever a LF representation is used to generate a phonological representation of a sentence, then the sentence appears as a consciously 'heard' sentence in inner speech.

Now we have some room for explicit, exact thoughts: they are the unconscious LF thoughts occurring 'behind the scenes', so that for each fragmentary sentence in conscious thought that our introspection reveals, there is a complete, explicitated unconscious thought, in LF, that 'backs it up'. The reasons why the complete thought does not surface to consciousness may well have to do with the additional computational load required to entertain a conscious thought. We can reduce the load by presenting to consciousness only the most conspicuous elements of the (unconscious) thought that we have in mind. Starting from those fragments, the unconscious machinery would be capable of reconstructing the corresponding complete and exact thoughts in NL.

9   Fodor (1998b, pp. 64–65) follows a similar, even if somewhat underdeveloped, line of reasoning. In his view, it might be the case that we are able to think in an ambiguity-free regimentation of English, roughly, in formulas of logical form (the level, on the other hand, favored by Carruthers (2002), see below), but this would not be the sort of linguistic structure that is given to introspection.

Yet, in taking this move, Carruthers would be reversing the direction of the evidence with which he initially motivated his view. His point was to draw conclusions about the nature of thought, starting from the data that introspection affords. We have produced reasons to doubt that those data offer genuine NL sentences. So the argument ought to lead us in the direction of denying the linguistic nature of other domains of thought. But the rationale of the previous paragraph was the opposite: the linguistic nature of unconscious thoughts was posited in order to save the linguistic nature of conscious thoughts.

Does Carruthers have an additional independent argument to support the view that some unconscious thoughts ought to be framed in NL? Carruthers (1996) offers two brief considerations in this regard. He claims that language is involved in all those thoughts whose constituents depend on language for their acquisition, and he regards it as a more economical and unifying hypothesis to treat conscious and unconscious thoughts as the same linguistic kind. With respect to the first point, it can be replied that it is one thing to be involved (even essentially) in the acquisition of a concept, and another rather different one to *constitute* such a concept. With respect to the second point, it needs only be said that it is as much an economical hypothesis to treat both conscious and unconscious thoughts as the same *nonlinguistic* kind. The only argument that seems to be left is that based on the evidences for the role of language in intermodular integration. As we said above, treating this position at length is beyond the scope of this paper. Yet, we want to end by arguing that even this line of evidence, as well as the introspectivist data that lend the initial momentum to Carruthers's position, can be explained without reaching the conclusion that language is a vehicle of thought.

7. A Cognitive Role for Language

We have argued that, even if language were potentially compositional and could, in principle, avoid the problem of semantic underdetermination, we do not have compelling evidence of the presence of non-underdetermined sentences in consciousness. And, given that thought cannot be underdetermined, we may conclude that natural language is not a vehicle of thought. At this point, one might want to consider a final way to escape this conclusion: perhaps we should never have conceded that thought is compositionally exact. After all, this hypothesis may be too close to the classical, sentential paradigm championed by Fodor, and this paradigm has been under heavy attack from different sides. Schiffer (1987, 1991), for one, has argued that neither NL nor the language of thought has a compositional semantics. Clark (1993), from a connectionist perspective, suggested that thought does not have the sort of context-independence that compositionality (classically understood) requires. Perhaps it is possible to develop a coherent defense of the cognitive view of language following one of these trails. We will not do so, first, because it would take us well beyond the limits of this paper and,

second, because it is doubtful that Carruthers would approve this move, concerned as he is with highly-structured thought.

What we are going to do instead, to close this paper, is to sketch an alternative reading of Carruthers's evidence, based on, though not identical with, Clark (1997, 1998). Drawing our attention to some habitual uses that people make of language, he has brought forward a cognitive role for language that does not turn it into an instrument of thought, but only into a facilitator (or, as he likes to say, a 'scaffolding') for thought.[10] In our view, this thesis can be sustained independently of his views on the nature of cognition as fundamentally a pattern-recognition process whose structure need not mirror any linguistic properties, and independently of his assent to the thesis that inner speech is constitutive of conscious propositional thought.

Quite often, in writing or speaking, one discovers oneself having thoughts that otherwise one would not have. By 'objectifying' the thoughts and contemplating them, be they written or spoken, one has a different access to them, which seems to provide a different stimulus. According to Clark, language's proper function is to communicate thoughts; however, we can give it a derived function given its capacity to objectify our thoughts, and thus convert it into an appreciable help in order to reflect, revise and, in general, facilitate our access to thought. In Clark's words, language has the capacity to give birth to a 'second-order cognitive dynamics', i.e. a return to, and revision of, our thoughts. This way, language becomes an external tool that we use in order to gain knowledge about our own mental life. Which, by the way, explains why there is a correlation between having a language and being capable of deliberation.

However, there is a more interesting sense in which language is a cognitive tool. Following a Vygotskyan trend, Clark points out that some external tools can be internalized. For instance, we make multiplications with the external help of pencil and paper, but we also do it by internalizing that external help and imagining numbers written on a piece of paper. Language, therefore, becomes something that can be perfectly considered as internal to the mind, and the cognitive uses that we make of it as an external tool become cognitive uses of part of our mental life.

Therefore, according to Clark, language is a tool whose proper function is communication but which, by virtue of that function, and thanks to our capacity to internalize some external tools, becomes an unbeatable cognitive dynamist. By internalizing language in this way, we create a special kind of mental object that facilitates having certain thoughts, allows us to fix our attention upon them and bring them back to memory in an easier way. And all this without having to 'go outside'. This seems to be a good account for the uses of language in thought listed above (thinking for writing, imagining a dialogue, rehearsing instructions . . .), which involve the contemplation of what is thought as if it were an object that

---

[10]   A similar role for language has been suggested by Jackendoff (1997), from assumptions having to do with his theory of consciousness. The moral of the story is pretty similar to the one told in this section, given that the cognitive role of language arises from its being at the interface between mind and world.

has to be perceived from several different points of view. In general, they are uses that are involved in hard cognitive tasks, thus it is no surprise that, as in the case of multiplying, we have to use an external tool. In addition, Clark's position allows that, as Carruthers defends, information arriving from different modules comes together by linguistic means. Yet, to fulfill this duty it is not necessary to regard language as a vehicle of thought: we can simply conclude that language provides a modality-neutral space where data incoming from different modality-tied modules can be brought under the same focus. The notion of integration need not be understood, as Carruthers contends, in terms of NL's providing the *syntax* necessary for intermodular integration. Rather, language may simply provide a means for thought to concentrate the attention on, and hence to combine by thought's own compositional means, those mental representations more relevant to the task.[11]

Admittedly, this sketchy view is highly speculative. It is not sure that we do not use language in other processes, even that we do not 'introspect' the use of language in other, perhaps simpler, cases, and maybe information integration takes place inside the language module. Yet the argument that we have been rehearsing here leads us to conclude that the information that the linguistic module ultimately produced should be subjected to interpretation by thought. Given that an account like Clark's is perfectly compatible with this interpretational demand for language, we claim that it does not face the same difficulties as Carruthers's position. Underdetermination problems suggest that NL can be at most a help, perhaps a necessary one, for thought. To sum up, we take it that, given the introspective data of a very fragmentary use of NL and the problem of underdetermination, the account that has it that we mainly 'talk to ourselves' in order to help us to think is a promising foundation for the whole issue of the cognitive use of language.

8. Conclusions

We have shown how semantic underdetermination, understood in Recanati's sense, prevents natural language from becoming a vehicle of thought. It has been conceded that, even if natural language were able to produce complete sentences (i.e. that SU were not an irredeemable feature of language), the linguistic items that we experience introspectively do not belong to this kind. That is, the NL sentences that we experience in cognition are indeed plagued with semantic underdetermi-nation. Hence, it cannot be the language that we use in thinking. However, we

---

[11] Carruthers himself characterizes his view in contrast with Clark's in the following sense: while the former deals with thought tokenings, the latter focuses on the process of thinking extended over time. (We thank a referee for calling our attention to this point). We do not think that this difference in focus is of great consequence for Carruthers's position. The role that a Clark-based view attributes to language can be also envisaged in a static way, i.e. linguistic items afford a type of anchoring for those elements of thought that one wants to highlight.

have tried to show that this does not mean that we do not use natural language in cognition. We most probably make use of language as an internalized external tool. We find this position not only intuitively appealing, but also the account (or sketch of an account) that behaves best when confronted with the data and arguments here developed.

First of all, it explains our introspective data. Like Carruthers's account, it takes them at face value, but unlike it, it explains what we take to be the real data (i.e. radically incomplete sentences). Second, it has no problem with the fact that sentences in natural language are usually underdetermined, for the role here ascribed to language does not require that its sentences are compositional and exact. There is no reason why an incomplete expression should not help us reflect about, or attend to, a thought, if we are able to interpret such an expression. We can do this provided that (a) we know well enough what that expression is intended to mean, and (b) the language we use in interpreting it is free of under-determination. As it is usually the case that we know what we mean by the expressions we utter to ourselves, and we do have a complete exact language of thought that we can use to interpret the pieces both of spoken and thought language, the phenomenon of SU poses no problem to the account endorsed here. Third, this metarepresentational hypothesis does not *require*, of course, that NL suffers from underdetermination. NL can be used as a help in the present sense whether or not it is explicit. However, if NL is, in effect, inexplicit, Clark's or a similar hypothesis becomes more plausible inasmuch as it makes an alternative view—Carruthers's—less reasonable. Finally, this account can also explain why we use natural language at all, given that we have a language of thought different from any natural language and free of its problems. To repeat, we use natural language as a tool, in the same way that we use imagined written numbers in order to make a multiplication. We use natural language to make our mental life simpler. The account here sketched can explain why we would 'speak to ourselves', i.e. why we speak to someone who knows what we are going to say. The reason is, again, that we tell things to ourselves in order to help us to think.

Departamento de Filosofía
Universidad de Valladolid,
Spain

Departamento de Filosofía
Universidad de Granada,
Spain

References

Bach, K. 1994: Conversational impliciture. *Mind & Language*, 9, 124–62.
Bickerton, D. 1995: *Language and Human Behavior*. Washington: University of Washington Press.

Carruthers, P. 1996: *Language, Thought, and Consciousness: An Essay in Philosophical Psychology*. Cambridge: Cambridge University Press.

Carruthers, P. and Boucher, J. (eds) 1998: *Language and Thought: Interdisciplinary Themes*. Cambridge: Cambridge University Press.

Carruthers, P. 2002: The cognitive functions of language. *Behavioral and Brain Sciences*, 25, 657–726

Clark, A. 1993: *Associative Engines*. Cambridge, MA: MIT Press.

Clark, A. 1997: *Being There*. Cambridge MA: MIT Press.

Clark, A. 1998: Magic words: how language augments human computation. In Carruthers and Boucher, pp. 162–83.

Dennett, D. C. 1991: *Consciousness Explained*. Boston: Little, Brown.

Fodor, J. 1975: *The Language of Thought*. Cambridge MA: Harvard University Press.

Fodor, J. 1987: *Psychosemantics*. Cambridge MA: MIT Press.

Fodor, J. 1998a: *Concepts*. New York: Oxford University Press.

Fodor, J. 1998b: Do we think in Mentalese? Remarks on some arguments of Peter Carruthers. In J. Fodor, *In Critical Condition*. Cambridge, MA: MIT Press.

Fodor, J. 2001: Language, thought and compositionality. *Mind & Language*, 16, 1–15.

Fodor, J. and Pylyshyn, Z. 1988: Connectionism and cognitive architecture: a critical analysis. *Cognition*, 28, 3–71.

Frawley, W. 1997: *Vygotsky and Cognitive Science*. Cambridge, MA: Harvard University Press.

Jackendoff, R. 1997: *The Architecture of the Language Faculty*. Cambridge, MA: MIT Press.

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.

McDermott, D. 1981: Artificial intelligence meets natural stupidity. In J. Haugeland (ed.) *Mind Design*. Cambridge, MA: MIT Press.

Mithen, S. 1995: *The Prehistory of Mind*. London: Thames and Hudson.

Pinker, S. 1994: *The Language Instinct*. New York: Morrow.

Recanati, F. 1995: The alleged priority of literal interpretation. *Cognitive Science*, 19, 207–32.

Recanati, F. 2001: What is said. *Synthese*, 128, 75–91.

Recanati, F. 2002: Does linguistic communication rest on inference?. *Mind & Language*, 17, 105–27.

Schiffer, S. 1987: *Remnants of Meaning*. Cambridge, MA: MIT Press.

Schiffer, S. 1991: Does Mentalese have a compositional semantics? In B. Loewer and G. Rey, (eds), *Meaning in Mind: Fodor and his Critics*. Oxford: Blackwell.

Spelke, E. and Tsivkin, S. 2001a: Initial knowledge and conceptual change: space and number. In M. Bowerman, M. and S. Levinson, (eds), *Language Acquisition and Conceptual Development*. Cambridge: Cambridge University Press.

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.

Stanley, J. 2002: Making it articulated. *Mind & Language*, 17, 149–69.

Stanley, J. and Szabo, Z. G. 2000: On quantifier domain restriction. *Mind & Language*, 15, 219–61.

Varley, R. 1998: Aphasic language, aphasic thought: an investigation of propositional thinking in an apropositional aphasic. In Carruthers and Boucher, pp. 128–45.

Vygotsky, L. S. 1934: *Myshlenie i Rech': Psykhologicheskie Issledovaniya*. Moscow and Leningrad: Gosudarstvennoe Sotsial-no-Ekonomicheskoe Izdatel'stvo. (*Thought and Language*. Cambridge, MA: MIT Press, 1962).

Waismann, F. 1951: Verifiability. In A. Flew (ed.), *Logic and Language*, 1st series, Oxford: Blackwell, pp. 17–44.