

Epistemology of Language

Edited by

ALEX BARBER

OXFORD
UNIVERSITY PRESS

Grasping Objects and Contents


Reinaldo Elugardo and Robert J. Stainton

I An Example

Imagine Andrew walks into a room, holds up a cigarette, and says ‘From France’ to Sylvia. We think it is obvious that, in this case, Andrew could easily convey a proposition. Let us agree that Andrew communicates, about the cigarette, that it is from France. A singular proposition. He can do this because, as will be obvious to both Andrew and Sylvia, what Andrew means clearly is not the property $\lambda x . \text{from-France}(x)$. How could he mean *that*? Thus, to treat him as co-operating, Sylvia must find a proposition meant, and the most obvious one is precisely this singular proposition. Supposing that what is communicated is this proposition, consider now what expression was produced. We maintain, though this is not entirely obvious, that Andrew did not utter a sentence. He did not, we maintain, utter the ordinary sentence ‘This cigarette is from France’. Nor did he utter any of: ‘This thing is from France’, ‘The item in my hand is from France’,

We are grateful to the Social Sciences and Humanities Research Council of Canada for a grant held jointly by the two authors. Among other things, this grant provided support for a Research Assistant, Jennifer Schellinck. We are also grateful to the University of Oklahoma Research Council for a grant that helped in the production of this paper. Thanks also to Alex Barber, Anne Bezuidenhout, Andrew Botterell, Robyn Carston, and Zenon Pylyshyn for comments on earlier drafts, and to the participants in the Epistemology of Language Conference at the University of Sheffield in 2000, where this paper was read. We must especially single out Emma Borg, who presented a very insightful and helpful commentary at that conference, as a result of which the paper has been improved in many ways. An earlier draft was presented to the Department of Philosophy at the University of Missouri (Columbia), and we are grateful for the helpful comments we received on that occasion. Finally, special thanks to Ernie Lepore for encouraging us to pursue what, to some, will seem an excessively idiosyncratic topic.

or even 'This is from France'. What he uttered was the bare prepositional phrase:¹



It also looks, at first glance, as if Andrew's interlocutor *recovered* a non-sentence as well. (Again, this is controversial. We return to it below.) Prima facie, what Sylvia heard was precisely what Andrew produced: a prepositional phrase. Yet she understood the singular proposition that Andrew meant.

Note too: the *meaning* of this PP is subpropositional. For instance, it clearly does not contribute a proposition as it appears in 'I brought this wine back from France'. Its meaning as it appears in this sentence is some kind of property, λx . from-France(x). And it would be absurd to suppose that it and every other phrase in the language has *two* meanings, this subpropositional one (which it contributes within sentences and other larger structures) and some cluster of fully propositional meanings (one of which it exhibits when it occurs alone).² Now, the fact that a subpropositional expression can be employed to communicate a proposition should not be that surprising. In fact, there are lots of cases in which what the speaker communicates goes well beyond the meaning of her words. This surely occurs in conversational implicature, and many would maintain that it equally occurs in metaphor, indirect speech acts, irony, etc. We strongly suspect that it occurs when quantificational domains are contextually restricted (for example, 'Everyone got drunk on Friday' is used to mean, not that every person in the world got drunk on Friday, but that everyone in some salient group did so). It occurs when speakers refer to things using expressions that do not, even in

¹ Throughout this paper we simplify syntactic trees, using notation that will be more familiar to philosophers. For instance, in this tree we ignore the P-bar level. And we treat 'France' as an NP, not as a DP. We presume that these simplifications are only expository.

² Why a *cluster* of propositional meanings, rather than just a single propositional meaning? Because a given phrase can be used to communicate a proposition of subject-predicate form, of quantificational form, an identity, etc. And these cannot arise from the contextualization of a single propositional meaning. For the arguments see Stanton (1995).

the context, denote those things (for example, someone says 'Your mother is very tall', referring thereby to your much older sister). And so on. So, we think it is plausible that the same happens here: the hearer understands the proposition which the speaker meant, even though the speaker's words do not, even in context, mean that proposition.

Thus, because of pragmatics, a person (e.g. Andrew) can successfully communicate a complete proposition by uttering something subsentential, with neither the syntax nor the semantics of a sentence. (In the case at hand, Andrew uttered a prepositional phrase of semantic type $\langle e, t \rangle$.) That said, though a pragmatic process plays a part in determining the proposition, it does not seem that Andrew merely *implicated* a proposition. He asserted one. Certainly he could not later say, accused of lying about the origin of the cigarette: 'Actually, I made no statement at all. Neither about the cigarette, nor about anything else. Sylvia just drew inappropriate conclusions.' This would radically misdescribe the case.

Given that an assertion was made, it is very tempting, upon first encountering such examples, to dismiss them as 'elliptical'. But, we maintain, what Andrew did cannot be explained by appeal to the sort of thing that contemporary syntax calls 'ellipsis'. Of course Andrew 'spoke elliptically', in some pre-theoretical sense—Andrew definitely communicated more than what his words meant. What we deny, however, is that one can explain away what Andrew did by assimilating it to the previously understood phenomenon of, for example, VP ellipsis. It is, surely, an empirical issue whether we are right about this. But this is not the place to lay out the evidence in full—instead, we sketch some sample evidence in the sequel.

2 Background Issues in Epistemology and Philosophy of Language

Much of this paper will be dedicated to the question of how exactly a hearer of subsentential speech—e.g. a hearer of 'From France', said on its own—manages to arrive at the proposition meant. This may look like a merely psychological issue. We think it is not. To see why these sorts of example might matter to philosophy, we want briefly to remind the reader of some initially plausible philosophical views, in epistemology and philosophy of

language. We shall not be exploring these views here—our aim, rather, is to gesture at the philosophical backdrop of our paper. The views are:

- Only expressions of natural language have logical form. Or anyway, anything else which is assigned a logical form gets it derivatively from a natural-language expression. Thus if there are propositions, and they have logical forms, they do so only because they are expressed by natural-language sentences that have logical forms. Similarly, if there are 'Mentalese' sentences, and they have logical forms, they too must get them derivatively.
- All effects of context, at least on what is strictly asserted, are traceable to elements of syntactic structure. Thus, once all explicitly indexical 'slots' are filled in, what is literally said by an agent is exhaustively determined by the conventional meaning of the (disambiguated) expression. Anything else which the speaker might have meant is not part of what is said.
- Issues about disambiguation and reference assignment to indexicals aside, knowledge of a language *L* is sufficient for interpreting utterances in *L*.
- The semantics/pragmatics boundary amounts to this: semantics assigns truth conditions to utterances, thereby yielding what is literally asserted, while pragmatics goes beyond this.
- Assertion is a conventional activity. What makes something an assertion is not one's intentions, but the conventions governing a certain public linguistic practice. Specifically, assertion is the act of producing a declarative sentence under conventionally specified circumstances.
- Only sentences can be used to make a move in the language game.
- Only sentences have meaning in isolation.
- Specifically, quantifier phrases (including definite descriptions) do not have meaning-relata. Instead, their meaning must be given metalinguistically, in terms of how they affect the meaning of complete sentences.
- Because sentences are the only things which have meaning fundamentally, word meanings (if such exist) must supervene on the total class of sentence meanings.
- Word meanings are underdetermined by the complete class of sentence meanings. Therefore, given the above, word meaning is indeterminate.

As we have argued elsewhere, all of these claims are *prima facie* in tension with the genuine existence of subsentential speech acts. We shall not rehearse the arguments here. Instead, we shall merely gesture at where the tensions might come from.

Let us start with logical form. Thoughts communicated subsententially, including the Andrew–Sylvia example above, typically have a full-blown logical form. The proposition about the cigarette can, for example, serve as a premiss in an argument. For instance, suppose Andrew had in previous days been debating with Sylvia (who is a renowned Francophobe) about whether anything really fabulous had recently been produced in France. But Sylvia knows and appreciates fine cigarettes. So, Andrew addresses his remark ‘From France’ to a person who, both discussants know, will recognize the inherent value in the displayed cigarette. Which cigarette is, let us agree, of quite recent vintage. In this instance, Sylvia draws inferences on the basis of the cigarette proposition, concluding that her claim, that nothing fabulous is being produced in France, is false. The crucial point is, if the thing-meant is to serve as a premiss in an inference, it must have a logical form. Yet the hearer apparently does *not* recognize this logical form derivatively—e.g. by recovering a sentence that has it. Rather, she assigns it a logical form fundamentally. If this is right, people can non-derivatively assign logical forms to things which are not expressions of natural language. (For rather more on this point, see Elugardo and Stainton 2001.)

Subsentential speech also suggests that understanding what a speaker said requires rather more than knowledge of language. It even requires more than knowing the disambiguated structure/content of the thing uttered, and the referents of any of that structure’s indexicals. For, returning to the example, what Andrew said (i.e. asserted, stated, claimed) was a proposition; but the meaning of the prepositional phrase which he uttered, even after disambiguation and fixing of reference for indexicals, is a property, not a proposition. So knowing what Andrew said—knowing, for example, the conditions under which it would be true or false—requires knowing more than the structure and meaning of the thing he uttered. Knowing language, even knowing contextualized and disambiguated language, is not, therefore, sufficient for interpretation. This ties in very directly with the issue of the determinants of what is said. Some philosophers—e.g. Paul Grice in some moods, Jason Stanley in all moods—maintain that to arrive at ‘what is said’,

it is sufficient to assign reference to all elements of the syntactic structure, and disambiguate. Anything further, which the speaker might have meant, cannot be 'said' but must instead be merely 'implicated'. (This minimalist view has recently been challenged by, among others, Carston 1988; 2002; Récanati 1989; Sperber and Wilson 1986; and Travis 1985. It is defended in Stanley 2000.) Assuming that the phrase 'what is said' is used here in the sense of what is asserted, stated, or claimed, this view seems also to be falsified by subsentential speech, as Stainton (1997a; forthcoming) argues.³ For, as noted, agents who speak subsententially do not merely *implicate* propositions. They assert them. Yet, as was just seen, the meaning of the disambiguated reference-assigned expression, in the case of subsentential speech, is (typically) subpropositional. So the content of what speakers assert often goes well beyond the meaning of the expression used, even after reference assignment to indexicals. In the same vein, subsentential speech suggests that there are effects of context on what is said with no syntactic counterpart. (Here again, see Stainton forthcoming and Elugardo and Stainton forthcoming, which respond to Stanley 2000. See also Clapp 2001 and forthcoming for a critique of Stanley.)

This result about the determinants of what is said clearly has implications, in turn, for the analysis of assertion and for debates about where to draw the semantics/pragmatics boundary. *Pace* Dummett (1973), for example, assertion cannot be analysed as the production of a declarative sentence in conventionally specified circumstances, because one can use subsentences to assert, and to make other 'moves in the language game'.⁴ Nor will it do to say, 'Semantics is about truth conditions'—so that semantics assigns truth conditions to utterances, thereby yielding what is literally asserted, while pragmatics goes beyond this. This will not do because if subsentential speech is a genuine phenomenon, pragmatic processes play a key role in determining the truth conditions of the utterance.

³ This caveat is important. Some authors, most notably Kent Bach (1994a; 1994b; and elsewhere) essentially *define* 'what is said' as the result of disambiguation and reference assignment. In which case, it is of course true that 'what is said' never goes beyond this. Still, Bach allows that what is asserted/stated/claimed may go well beyond 'what is said', so defined. Indeed, what he calls 'implicatures' count as contributions to what is asserted that go beyond 'what is said' (in his sense). So any disagreement between Bach's views and the stance taken in the text is (mostly) terminological.

⁴ For discussion see Stainton (1997b). An insightful reply can be found in Kenyon (1999).

Consider, finally, Frege's context principle, and its (supposed) implications. If only sentences have meaning in isolation, then subsentences do not. So, in particular, quantificational expressions cannot be assigned meaning-relata. Yet, as Stainton notes (1998a; 1998b), such phrases can be used, and understood, on their own. (See also Botterell forthcoming.) What, then, is being claimed, when it is said that they lack meaning in isolation? Must it not be the case that users *know* the meaning outside any sentence, and can employ that meaning without a sentential environment? Another implication: understood in certain ways, the context principle can be used to support word-meaning indeterminacy, on the grounds that word meaning, if it exists, must supervene on sentence meanings—this because only sentences *have* meaning non-derivatively. It is often added that, as a matter of fact, there are many possible lexical entries for each word consistent with the complete set of meaning-specifications for whole sentences. So word-meaning must be indeterminate, since it is left underdetermined by the reputed source of 'meaning facts', taken in its entirety. But if subsentences can be *used* *sō* freely, in speech, in precisely what sense do they lack meaning in isolation?⁵ And if non-sentences have meaning in isolation, why should their meaning have to supervene on sentence meanings? Why can they not have meaning non-derivatively? In sum, what the 'context principle' amounts to, what it entails, and whether it is even true, will all three have close ties with the phenomenon of non-sentence use. Or so it appears.

Granted, appearances could mislead. Not just with respect to the context principle, but with respect to all the foregoing doctrines. It is not obvious that the *prima facie* implications for epistemology and philosophy of language really will obtain, once all is said and done. Countermoves for explaining away the apparent tensions are certainly available. (See e.g. Barber forthcoming; Davis forthcoming; Kenyon 1999; forthcoming.) On the other hand, one equally cannot rule out *a priori* the relevance of subsentential speech to issues like indeterminacy, the domain of logical form, the

⁵ Maybe it is true that to give the meaning of a word/phrase, one must say what other meanings it could combine with to yield something truth-evaluable. This certainly seems a plausible methodological precept. If so, this would maintain the centrality of truth, or Fregean Thoughts, for lexical semantics. But it would not entail that *sentences* have any special place in natural language. And it certainly would not entail that the only supervenience base for word meanings was complete sentences. (For discussion see Stainton 2000.)

determinants of what-is-said, the nature of assertion, the context principle, etc. (Philosophy of language owes debts to empirical work on language, just as philosophy of biology owes empirical debts to biology.) We think that the mere possibility that these doctrines might be in trouble ought to be sufficient motivation for philosophers to consider the issue of subsentential speech seriously.

3 Our Questions and Aims

We began by introducing an example of subsentential speech. We then highlighted, in the previous section, *prima facie* tensions between subsentential speech and a wide range of widely held views in epistemology and philosophy of language. As noted, we shall not explore those doctrines further here. Instead, our central question in this paper is

(Q1) ‘How do hearers interpret subsentential speech of this sort?’

We do not give anything like a definitive answer to (Q1). What we do instead is sketch the salient options, and provide arguments—some philosophical and some more empirical—to the effect that certain of these options are implausible.

One reason for addressing (Q1) is that we find it inherently interesting. In part it is interesting because it affords a case study of the fascinating phenomenon of *cross-modal integration*, here of information got from language with information got from visual perception. This is

(Q2) ‘How is perception-based information seamlessly combined with communication-based information?’

What is rather novel about the present paper is how we approach this larger question: by looking at a seldom-studied variety of linguistic communication, namely less-than-sentential speech. Another quite different motivation is this. Defenders of the above philosophical claims can respond to the *prima facie* tension in two ways. On the one hand, they can insist that the tension is only apparent. This will not be our concern here. On the other hand, they can deny that the phenomenon of subsentential speech is genuine. It is this latter strategy we resist here, by arguing that the options for treating

the phenomenon as not genuine are implausible. In contrast, there is an empirically open option, the one we tentatively endorse, according to which the non-sentential speech phenomenon is genuine. This leads to our third question in this paper:

(Q3) 'In the light of the answer to (Q1), is it plausible that the phenomenon of subsentential speech is genuine, or is it more plausible that it is merely apparent?'

Having noted what our aims and questions are, let us also clarify what they are not. To repeat, for the most part, we shall not argue that apparent tensions between familiar philosophical doctrines and non-sentential speech are real; in the light of this, we equally shall not argue that these philosophical doctrines are actually false.⁶ The exception to this rule is a lengthy discussion of the role of 'inner speech' in occurrent thoughts. Here we do question the truth of certain philosophical doctrines. The doctrines in question are:

- Having an occurrent thought is a matter of having a sentence of natural language pass through one's mind.
- Having a genuine thought about an external object requires that some natural-language singular term, which refers to that object, pass through the mind.

Again, we take these to be in tension with the existence of the sort of example sketched at the outset.⁷ We shall argue that this tension is real. And we shall conclude that these two doctrines are false. To anticipate:

⁶ As noted, we have done both things elsewhere. Some of the relevant papers include: Elugardo and Stainton (2001; forthcoming, *a*); Stainton (1995, 1997*a*; 1997*b*; 1998*a*; 1998*b*; 2000; forthcoming).

⁷ We distinguish this version of the inner-speech view of thought from the weaker analogical version. According to the latter, thoughts are inner episodes that are psychologically type-individuated by their functional/conceptual roles in an internal system of representation; the roles they play are said to be formally analogous (in the appropriate ways) to the functional/conceptual roles played by sentences in public natural languages. Public natural-language sentences are also said to be conceptually prior to thoughts in the order of explanation. The idea is that the intentionality of thoughts is to be explained on the model of the semantic referential properties of natural-language sentences, which in turn get explained in terms of their conventional use in communicational contexts, rather than the other way around. Sellars (1963) is the classic defence of the analogical account of thoughts. Some recent, Sellarsian-inspired accounts of thoughts can be found

this hard-line view conflicts with the existence of genuine subsentential communication because hearers can understand complete propositions—e.g. the singular proposition about the cigarette, to the effect that it is from France—even when what passes through their mind is an expression that merely encodes a property. Thus there is more to the thought grasped than what appears in the supposed vehicle for it. Specifically, and interestingly, it seems that something *in the environment* is part of the thought grasped: in the case at hand, the cigarette itself is part of the thought meant by Andrew and understood by Sylvia. And this external object gets to be part of the thought grasped even though, or so it appears, no natural-language expression denoting this object passes through the mind of the agent who grasps the thought. Finally, continuing with what we shall *not* do, we shall not try to answer (Q1) with respect to all cases of subsentential speech. We are trying to take only the very first steps here, and our example is highly simplified as a result. Specifically, we shall be addressing only cases in which (1) the speaker produces a predicate and (2) the object talked about is visually attended to. As will emerge at the end of the paper, there are lots of cases of subsentential speech that are more complex than this. Thus, an enormous amount of work will remain to be done even if we get things right with respect to visually grasping an object and making a claim about it subsententially.

4 The Options

As it stands, we can conceive of only four approaches to (Q1), including especially the ‘visual grasping of the object’ in the environment that is being talked about. The first divide between approaches is what we call ‘S-representationalist’ vs. ‘Non-S-representationalist’ approaches. By ‘non-S-representationalism’, we mean the view that beliefs, thoughts, and the like do not have compositional syntactic structure. In particular, non-S-representationalists reject what William Lycan calls ‘Forthright Sententialism’: the view that all (occurrent) tokens of beliefs, thoughts, etc. are inner

in Gauker (1997) and Lycan (1988). None of the arguments that we raise in this paper against the ‘hard line’ version of the inner-speech view affects the analogical version. Considerations against the analogical theorist’s conceptual-priority claim are, however, raised in Elugardo and Stanton (2001).

brain states that exhibit syntactic structure, have truth conditions in virtue of referential relations between their 'syntactic elements and things in the world', and further, 'stand in causal relations to each other of the sort that constitutes inference' (Lycan 1990: 149). In short, S-representationalists embrace some version of the 'Language of Thought' Hypothesis: the view that thinking occurs in an internal language-like system of representation; non-S-representationalists, in contrast, reject any minimally interesting version of the LOT Hypothesis.

Within S-representationalism, there are two main suboptions. On the first, the inner thought-sentences are sentences of the speaker's native language: for instance, sentences of English. Such sentences could in principle include either ordinary sentences, or special 'elliptical sentences'. We pause to stress: if either 'natural-language sentence' suboption is true, then the appearance that people understand subsentential speech without processing a complete sentence must be a mere appearance. Thus, recalling Sylvia and Andrew, this view would entail that Sylvia recovered either an ordinary sentence (e.g. 'This is from France') or some elliptical sentence. Either way, it would not be true that she understood a thought that went beyond what was encoded in the words passing through her mind.

The other suboption within S-representationalism is that, though there are internal representations with both syntactic structure and compositional semantics, this 'language' is not a spoken language. Instead, it is some kind of 'Mentalese'. Within this latter option, the question arises as to whether the items which denote external things, thereby allowing the agent to 'grasp them' mentally, are exclusively descriptive or whether they (or anyway, some of them) might be indexicals. These two 'Mentalese' suboptions, in contrast to the prior 'natural-language sentence' ones, would yield that our phenomenon is genuine. For example, the only natural-language thing which passes through Sylvia's mind is a PP, and it is not of semantic type $\langle t \rangle$ even in context. The thing which is of semantic type $\langle t \rangle$, and which purportedly passes through her mind, is not a natural-language expression—it is a Mentalese sentence (containing an indexical, or not, depending on the approach proposed).

We shall explain each of these suboptions more fully in the sequel. For now, we summarize the lay of the land in Figure 8.1.




FIG. 8.1. The Options

4.1 *The non-S-representationalist option, (a)*

We introduce the non-S-representational view mostly to put it aside. By way of explaining it, we begin with the following rather rough-and-ready way of thinking about the original example. The hearer, goes this story, understood the prepositional phrase, and grasped the property (i.e. the propositional function) $\lambda x . \text{from-France}(x)$ for which it stands; she also noticed the cigarette that Andrew was holding up. She input the cigarette into the propositional function, and came thereby to grasp a singular proposition, about the cigarette, to the effect that it is from France. It was this process, and not the recovery of a sentential representation, *either in English or Mentalese*, that actually occurred.

Now, put this baldly, the approach seems to commit a serious use-mention confusion. 'She recovered the propositional function; then she input the object into the propositional function.' What can this mean? She did not physically pick up the cigarette, and put it into some kind of grinder, which then yielded a singular proposition at the far end. All of the 'recovering' and 'applying' occurred in her head. The underlying complaint here is that whenever a mental operation occurs—grasping an object, forming a thought about it, recognizing a property—some *representing* goes on. If that is right, then, one might think, something like the following must happen in non-sentential speech situations: the hearer tokens a representation of the object (i.e. a singular term) and also a representation of the property in question (i.e. a predicate); she combines these into a sentence; and she thus arrives at the complete thought. Hence, one might suppose, option (a) is a total non-starter.

This criticism of the general approach is far too quick, however. Of course

something must be happening to the agent, as she notices, considers, and so forth; and, if we are physicalists, that something must be broadly speaking material. Moreover, and equally obviously, the process of understanding non-sentential speech does not involve putting some external object into some sort of grinder. But this alone does not entail that this 'something' which happened in the head was a tokening of (a series of) quasi-linguistic symbols. Here are some reasonable alternatives. What happened to the agent was that her informational state changed—thereby altering her 'representing state', but without necessarily tokening any specific representation. Or, in a quite similar vein, what happened to the agent was that different counterfactual conditionals came to be true of her—counterfactuals about what she would do or say, for instance; again, without some sentence being 'written' anywhere in her head. Or, what happened was that her 'neural nets' got reconfigured in complex ways, once again without any tokening of a sentence-like representation. All of these are live options in the philosophy of mind. And it is at least possible that some or all of them could be used to interpret talk of 'grasping the object' and 'putting it into the function', without appealing to *internal sentence-like things*.⁸

Going 'non-S-representational' in this sort of way might be an option for explaining how less-than-sentential communication functions. But we shall not consider it further here. Suffice it to note that we have our reasons for being sceptical of any non-S-representationalist view. Our main reason is this: the central aim of our paper is to explain what might be going on in subsentential speech when we perceptually grasp an object and *then* combine this perceptual information with independently derived information, got from a different (i.e. linguistic) source, to form eventually a complete thought about the perceived object. To make sense of this, we need to be able to separate object-grasping from content-grasping, as distinct and

⁸ It should be noted that we include, in the category of 'non-S-representationalist approaches', such philosophical views as those of Robert Stalnaker, who defends an information-theoretic account of singular thoughts; cf. Stalnaker (1990). It would equally include those connectionists who insist that neural nets are (holistic) representations. Clearly these theorists do not reject mental representations of all sorts. But they equally do not think of representing as a matter of tokening inner *sentences*. The other camp, within non-S-representational approaches, deny the existence of mental representations of any variety. Connectionists of an eliminativist stripe would fall here, as might certain philosophers inspired by Davidson or McDowell.

isolatable processes. We suspect that a theory that posits mental representations that have a constituent structure has the best chance of explaining this type of phenomenon. By contrast, non-S-representationalist theories explicitly deny that our mental states have any combinatorial constituent structure. The minimal mental state they allow is, we take it, the comparatively holistic one of grasping a complete thought. Thus, by our lights, non-S-representationalists will be unlikely to be able to give a satisfactory account of subsentential talk. (We hope to take all of this up at length in a future paper.) That said, we do not argue in this paper that the only account of singular thought which can handle subsentential speech is one which posits the manipulation of internal, syntactically structured representations. That would take us too far afield, given the number of such views that have been defended in the literature. Rather, our strategy is to assume, for the sake of argument, that (what we are here calling) S-representationalism is correct and see what *kind* of S-representationalism best explains the phenomena of non-sentential speech, given certain empirical data about perception. (To anticipate: we shall argue that only a form of S-representationalism that posits a language of thought that is not a natural language has the best chance of explaining the phenomena. In which case, the phenomenon is genuine, and not merely apparent. In which case, those who endorse the philosophical doctrines noted at the outset have some work to do.)

4.2 *The natural-language-sentence options, (b) and (c)*

We begin our detailed discussions with options (b) and (c), within the natural-language-sentence approach. According to this view, there is a language of thought, namely, the natural language that a speaker/hearer actually uses in writing or in speaking. According to the view we have in mind (see note 7) all thought-tokens are internal speech-tokens of natural-language sentences. Given this, here is what a proponent of the view would probably say about (Q1) and (Q2). Since all thought is inner speech, whenever a person notices an object, she must subvocally token a natural-language singular term (e.g. 'that') which, in the context, stands for the object noticed; and when this person considers whether the object has a certain property, she tokens a natural-language sentence—containing both a singular term which stands for the object, and a predicate which encodes

the property (e.g. 'That is red'). (When it is stated this baldly, one might think that no one really endorses this line, but some philosophers actually do endorse it or are at least committed to it.⁹) Adopting this approach, perceptual information must be combined with communication-based information by having all the information in question encoded in natural language. In which case, the hearer of (apparently) non-sentential speech must have somehow tokened a complete natural-language sentence, which then completely encodes what the speaker meant. Applying this to the example noted above, such a theorist would have to say something like this: Andrew's hearer must have recovered some sentence, in her native tongue, which meant (something like) *This is from France*. Maybe she recovered the ordinary sentence 'This is from France', reconstructing it on the basis of the less-than-sentential words Andrew produced. Maybe she recovered some special 'elliptical sentence', which Andrew also uttered. But she definitely recovered some natural language sentence or other.

The question is, are such views really plausible? Let us take the options in reverse order, starting with (c). The following considerations should at least make our claim that no 'elliptical sentence' was uttered/recovered plausible enough for present purposes. (This matter is dealt with at much greater length in Stainton 1997a and forthcoming. See also Barton 1990; Carston 2002; Clapp 2001; forthcoming; Elugardo and Stainton, forthcoming a; and Morgan 1989.) Here is the first anti-ellipsis consideration. Syntactic ellipsis is a matter of grammatically necessary syntactic material 'being left out', but such that the omission does not matter because the very material which is omitted is present elsewhere in the discourse. Classic cases of syntactic ellipsis include VP ellipsis and 'sluicing', illustrated below:

- (2) Joan will buy a car but Laxmi won't [_{VP}].
- (3) Joan will buy a car. I wonder why [_S].

These sentences are bona fide cases—more than that, they are paradigm cases—of the phenomenon of ellipsis. (The first is VP ellipsis. The second, in which an S is elided, is the example of 'sluicing'.) Now, if what appeared to be subsentential speech were really just like this, then one could explain what was occurring between Andrew and Sylvia, for example, without

⁹ The inner-speech view of thought is defended in Carruthers (1996), in Devitt and Sterelny (1999), and in Ludlow (1999).

granting thoughts *sans* natural-language vehicles. On the other hand, if what one means by 'ellipsis' is *not* this, then 'ellipsis' is merely a comforting label, in lieu of an account. So, to give the hypothesis teeth, and explanatory power, we shall take it that the proponent of 'elliptical sentences' wants to assimilate the phenomenon to this already understood kind of case.

That said, notice two key features of the examples. To start with, in (2) and (3) the omitted material ('buy a car' and 'Joan will buy a car', respectively) is spoken elsewhere, just prior to the ellipsis. And this feature is shared by elliptical expressions in general.¹⁰ Syntactic ellipsis, in short, requires prior linguistic material. But then, to anticipate, what Andrew produced was not an elliptical sentence, because in the situation described there was no (immediately) prior linguistic material.

To see *why* prior material is required, we briefly introduce one standard view of syntactic ellipsis. It is by no means the only account, but it is representative enough to give the feel of the thing. According to a theory first proposed by Edwin Williams (1977), syntactic ellipsis works as follows. Covert anaphors, which Williams dubbed 'deltas', are base-generated within elliptical sentences. That is, ellipsis is not transformationally introduced: put metaphorically, elliptical structures are present already at the level at which the sentence parts are first put together. Deltas are, then, very like the null-subject pronouns one finds in languages like Italian and Spanish, that give rise to sentences like:

- (4) [pro] está. comiendo.
 pro be_{3rd pers. sing., present} eat_{present-continuous}.

¹⁰ The qualifier 'in general' is necessary because of certain special cases noted in Hankamer and Sag (1976). These are cases in which something is sufficiently salient in the environment for VP ellipsis to be comprehensible, though it still sounds awkward. For instance, while looking at a man standing on a ledge, someone could say 'He won't', meaning that he will not jump. Such an utterance is definitely interpretable. However, the use of words and phrases is neither as constrained as the Hankamer-Sag cases, nor does it sound in any way awkward. See Stainton (forthcoming) for extended discussion. We should also note special cases of 'frozen ellipsis', in which a set phrase has been created from what was originally an ellipsis construction. An example would be 'May I?', used to request a cigarette. If the aim of this paper really were to establish, once and for all, that a sentence is not used in the contested cases, a full discussion of both sorts of case would be required. Since, however, we are here trying only to make our anti-ellipsis claim plausible, we shall set both aside.

'[He/she/it] is eating.'

In such elliptical sentences, deltas occupy the places where, pre-theoretically speaking, material is 'left out'. They have no pronunciation, and no meaning *per se*, but they get linked to prior syntactic material at the point of interpretation. More concretely, Williams (1977) assigns the following sort of structure to (2):

(5) Joan will buy a car but Laxmi won't [_{VP} [_V Δ_1][_{NP} Δ_2]]

In context (and roughly speaking), Δ_1 is co-indexed with 'buy' and Δ_2 is co-indexed with 'a car'. This structure yields the proper interpretation for the second conjunct of (2), namely that Laxmi won't *buy a car*. Moreover, the structure accounts for why (2) is not pronounced the way 'Joan will buy a car but Laxmi won't buy a car' is: precisely because deltas have no pronunciation.

Suppose that William's story is right. Now, look again at Andrew's utterance. He walks into the room, and the first thing he says is 'From France'. Clearly there is nothing for any purported covert anaphors in 'From France' to be linked to in this situation. There is lots of physical context. But there is not the requisite linguistic context. There is, therefore, no way within those countenanced by the theory for the supposed deltas to be interpreted—because there is no prior material to which the deltas are linked. The hearer in this case cannot, then, link up the 'empty slot' with these prior words, because there are no prior words. (Put another way, if what Andrew produced really were an elliptical expression, in the sense familiar to syntacticians, it ought to sound precisely as odd as 'She doesn't', produced on its own. But 'From France' is not the least bit odd in this situation.) We conclude that what Andrew did looks to be importantly unlike syntactic ellipsis, as encountered in (2) and (3).

It will help make our point if we distinguish between (1) the agent as a whole making an all-things-considered inference, in principle using all the information at her disposal, about what sentence the speaker might have had in mind and (2) an automatic, formal operation of syntactic reconstruction, performed subpersonally within an informationally encapsulated module. VP ellipsis is a process of the latter sort: the language faculty does not 'guesstimate', on the basis of all evidence in principle available to the agent;

rather, it blindly churns out a structure, based solely on the representations input into it. Being automatic and encapsulated in this way, all the syntactic material required for the grammatical operation must be available to the linguistic module: where, basically, that required material consists in, first, the fragment actually spoken plus, second, *the immediately prior discourse*. That is why prior discourse is essential. Moreover, the latter must contain an expression identical (at the level of logical form) with the element omitted, so that the ellipsis site may be linked to it. So not just any prior discourse will do. (For the arguments see Sag 1976. The point is applied to subsentential speech in Stainton forthcoming, and Elugardo and Stainton forthcoming, a.) In sum, it is because ellipsis of this sort is a strictly grammatical recovery process that prior *spoken* material is required for ellipsis to be grammatically licensed. Without that material, no narrowly syntactic process could recover the 'missing' elements.

Our point is that the process of subsentential comprehension is not like this. At best it is the 'whole agent'—i.e. the agent in principle using all the information available to her, both linguistic and non-linguistic—who figures out what sentence the speaker had in mind. (We discuss this latter possibility below, under the rubric of option (b). This is not to say that the whole agent cannot interpret a person who employs a genuinely elliptical expression without appropriate prior discourse. *Whole agents* can understand discourse-initial ellipsis constructions, using all the information at their disposal. But the things understood sound awkward indeed.)

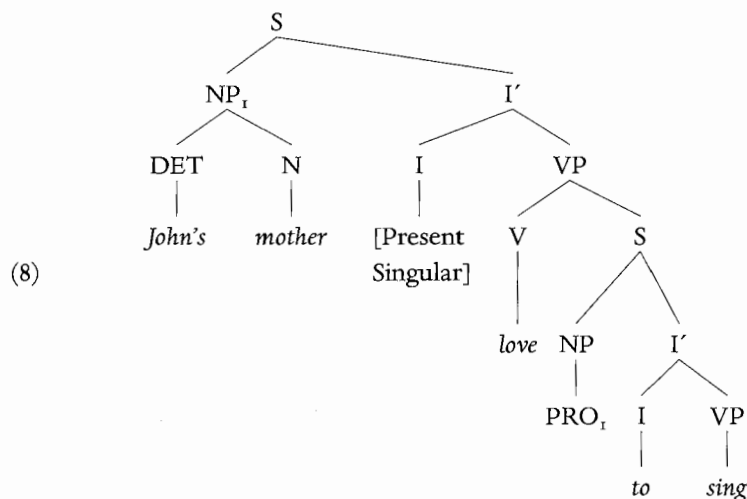
The foregoing is a pretty theory-internal bit of reasoning. Here is an argument which is more data-driven. Accounts of ellipsis agree that, as a matter of fact, genuine syntactic ellipsis, of the sort exemplified by VP ellipsis and sluicing, involves the omission of a syntactic constituent: a VP, an S, or some such. Absolutely no account of syntactic ellipsis countenances free deletion (and with no linguistic context) of whatever material happens to be necessary to leave a free-standing phrase. Thus, for example, (6) is ill-formed, as an elliptical version of (7):

(6) John's sing.

(7) John's mother loves to sing.

This is so even if it might be perfectly plain what is being talked about. (Thus Andrew might say 'Raymond's mother loves to dance'. If a non-native

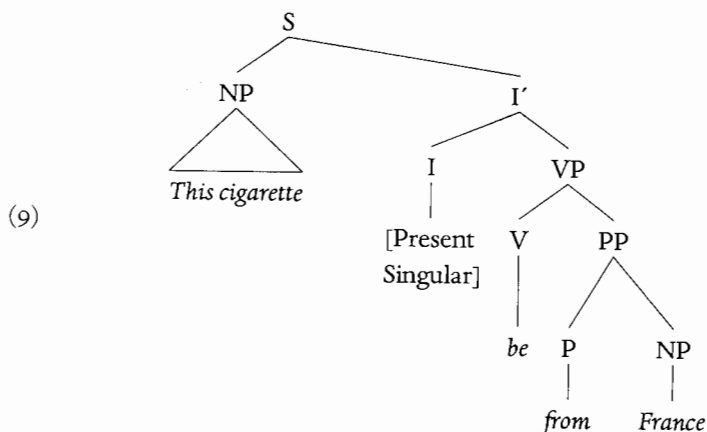
speaker responded with (6), we might well understand her as meaning what (7) says. But this is not syntactic ellipsis in the requisite sense: it is the whole agent forming an all-things-considered hypothesis about what the speaker 'really meant'.¹¹ One reason why (6) is ill-formed is that 'mother loves to' simply is not a syntactic constituent of 'John's mother loves to sing'.¹¹ The (simplified) tree for (7) is (8), and 'mother loves to' clearly does not fall under a single node in this tree.



Hence it is not possible to delete this non-constituent, so as to arrive at 'John's sing'.

That said, consider what would have to be omitted to derive 'From France' from its more obvious sources. Suppose the source sentence was 'This cigarette is from France'. The (simplified) tree structure for that sentence is given in (9).

¹¹ Again, there are complications even here. Though ellipsis of the VP sort requires that a constituent be omitted, *gapping* does not. Thus 'John's mother loves to sing, Jill's to party' is reasonably good, grammatically speaking. But what is omitted is 'mother loves', clearly not a constituent. However, whereas it might be plausible to suppose that Andrew performed something akin to VP ellipsis, it is not at all plausible that he produced a gapped expression. The reason is that gapping by definition involves leaving an empty 'gap' between two expressions. Clearly, however, if Andrew left anything out, it was not material in the middle of a sentence. For more on gapping and ellipsis see Berman and Hestvik (1992); Chao (1988); and Lappin and Benmamoun (1999).



The deleted material would have to be everything but 'from France'. But, patently, this is not a syntactic constituent. That is, it does not fall under one node. So, what would have to be deleted here would be a non-constituent, violating the generalization.

This is by no means a knock-down argument against treating the Andrew case as VP-style ellipsis: on the one hand, the generalization might be mistaken—indeed, one could take Andrew's utterance to be evidence against it. On the other hand, there might be a way of diagramming an appropriate source sentence such that what gets deleted really is a constituent. However, it certainly does not look to be syntactic ellipsis of the VP ellipsis sort, as this is currently understood. And that is sufficient for present purposes, since it therefore will not do to dismiss the question of how non-sentential speech gets interpreted by saying: 'Oh, we know how that works. It's just like VP ellipsis'. And, as noted, to say 'Well, it's not that kind of ellipsis. But it is ellipsis none the less' is to put a familiar label on the phenomenon, while refusing actually to explain it. A concrete account is still required.¹²

¹² We should note some examples which seem to pull the other way, i.e. in favour of an ellipsis account. As Peter Ludlow noted in discussion, one can use expressions like 'All in the garden' (to assert, for example, that all the cousins are in the garden) or 'By John' (to assert, for example, that a salient work of art was painted by John) in isolation. But these seem to be derived via sentence-level transformations, which suggests that the things uttered must actually be remnants of transformed sentences. See Ludlow (forthcoming) for extended discussion. (A similar argument was given independently in Morgan 1973, in favour of an ellipsis approach to what he termed 'fragments'. But see Morgan 1989 for

So much for option (c), which appeals to syntactic elliptical expressions. Let us now turn to option (b). We think the idea that to understand less-than-sentential speech one must recover an *ordinary* natural-language expression that picks out the element supplied by the environment is no more plausible than the idea that whenever someone notices an object, she tokens a natural-language expression that refers to it. It is highly implausible to suppose that, when someone looks at her desk and sees the objects on it, recognizing their features, a constant flurry of English sentences runs through her head. But then why suppose that when one notices an object *being discussed*, and considers its properties, one must token a singular term in English that refers to it? Certainly if we ask the hearer, 'Which complete sentence did you understand?', she will be at a loss to answer. At least sometimes. What makes one so sure that, in every such case, there will be the precise sentence used by the hearer? We say: nothing but a priori ideology. (Compare the case where Andrew says, 'This cigarette is from France', and we ask the hearer what sentence she understood.)

The argument against option (b) just presented presupposes the following principle: a hearer (i.e. the whole agent) recovers a natural-language sentence from an utterance of a subsentential linguistic item only if the hearer knows precisely which sentence she recovered. One might reject that principle on the grounds that the hearer could conceivably recover the sentence 'off-line', i.e. unconsciously; in which case, the precise natural-language sentence is potentially unavailable to the hearer's conscious mind.

We contend that this alternative 'off-line' picture, proposed on behalf of option (b), is equally unlikely. To see why, consider two accounts of

some rethinking.) Focusing on the first example, roughly speaking, it would be treated as having been derived as follows:

[[All the cousins] are in the garden] → [t_i the cousins are [all_i in the garden]] → [all in the garden]

But if this is how 'All in the garden' was derived, then ellipsis must have been performed on the (derived) complete sentence [t_i the cousins are [all_i in the garden]]. This seems the only way to generate this fragment. (Similarly for the case of 'By John'.) Nor will it do to say that *this* is genuinely a case of ellipsis, while 'From France' *et al.* are not, because 'All in the garden' can occur discourse-initially, and would have to be derived by deletion of a non-constituent. So if these are cases of syntactic ellipsis, the tests given in the text are not actually diagnostic. Such examples are discussed critically and in detail in Stainton (forthcoming). We ignore them here.

how the unconscious sentence would be constructed. On the first story, it is the language module which outputs the whole sentential structure. On the second, the language module outputs a subsentential string, and a complete sentence is built up using the Central System (in the sense of Fodor 1983), and all information available to the agent. Let us begin with the first, noting that the so-called 'recovered sentence' is not *completely* unavailable to consciousness.¹³ In our example, the hearer is consciously aware of the linguistic item that was pronounced in the speaker's utterance. Specifically, she is aware of the prepositional phrase 'From France'. So, if it is the language module which is outputting the complete sentence, certain of its results are available to consciousness, but certain others are not. Hence, if the inner-speech view is true, then the hearer is consciously aware of only part of the sentence that her language module outputs. Surely it would be odd for the language module to output a whole sentence only part of which is accessible to consciousness. But that is precisely what must happen if this first inner-speech theory is correct. Consider, then, the second story. Could it be that the language module outputs a subsentential structure, here a PP, such that it is the Central System's task to use all available sources to construct a natural-language sentence, given the subsentential structure as its input? This has a certain initial plausibility. Notice, for instance, that in the case at hand there is insufficient information within the language module to output a whole natural-language sentence, precisely because there is no syntactic antecedent. So the second story has that much going for it: it allows in contextual information not in the language module. But it will not work either. The activities of the Central System are typically accessible to consciousness. Thus, if this second inner-speech view were right, then the whole recovered natural-language sentence should, barring some argument, be available to consciousness, in line with the principle stated in the preceding objection to our argument. Neither proposal, then,

¹³ We also think that the reply goes against the spirit of the inner-speech view. One of the primary motivations for the inner-speech account is behaviourism. The idea is that thinking must be observable and public, at least in principle: it has to be either actual self-talk or dispositions to utter/accept natural-language sentences. Otherwise, we would have to relegate thoughts and the activity of thinking to the ghostly realm of covert, private, Cartesian mental objects and mental processes. Or so it was believed. Postulating a level of wholly unconscious mental processing is thus antithetical to certain aims of classical inner-speech theorists.

will account for why *part*—and *ex hypothesi*, only *part*—of the ‘recovered sentence’ is unconscious.

Although we think the above arguments are effective against the ‘ordinary sentence’ approach to subsentential communication, some might think otherwise. Also, some might think that the idea of having English sentences running through one’s head, as one attends to an array of objects, is no more implausible than the idea of having Mentalese sentences running through one’s head. (And we ourselves will endorse the latter idea.) We would like to offer, therefore, two other arguments against option (b).

First, it appears that a subject who clearly is cognitively deficient in natural-language understanding/processing *vis-à-vis complete sentences* can still understand non-sentential speech. Here are two considerations. First, presumably Helen Keller was already forming thoughts some time before she learnt how to sign full sentences. (Inner-speech theorists will probably deny that; in which case, they must explain how Keller was able to learn and understand sign language even though she lacked natural-language competence.) Suppose, then, that her teacher placed Keller’s hand in a bowl of water when she was teaching her the sign for ‘wet’. At some point during the learning process, Keller was able to understand her teacher as saying *that is wet* when the teacher made the sign ‘wet’. But it is doubtful that Keller tokened a natural-language sentence in her head that meant that: she did not yet have any at her disposal. Second, if the inner-speech theory is true, then it follows that as a person loses his competence to form and understand complete sentences, he also loses (in varying degree) his ability to think and reason. But, though typically true, this conclusion is not obviously correct across the board, given recent findings about language-specific deficits. For our purposes, the issue is whether a patient could still understand non-sentential assertions/questions/commands, etc., having lost the ability to use/understand complete sentences. If she could, then it is not plausible that, in the normal circumstance, someone who understands subsentential speech is, even unconsciously, deploying a complete sentence. Here is why: if understanding an (apparently) subsentential utterance demands constructing (and then understanding) a complete sentence, then someone who cannot construct a sentence should not be able to understand subsentential speech. The two capacities should both suffer equally, across the board. On the contrary, however, empirical studies appear to show that cer-

tain pathological individuals understand non-sentential speech even when their ability to construct complete natural-language sentences is severely impaired. That is, though they have lost many of their grammatical abilities with respect to natural-language sentences, including even the ability to distinguish ungrammatical from grammatical sentences, their comprehension of isolated lexical items (or sometimes even phrases) remains comparatively good. For example, the patient discussed by Chatterjee *et al.* (1995) had flawless comprehension of single words, and, as demonstrated by his response to pictures, could still understand concepts like agent of an action, recipient of an action, etc. But he could not encode this sort of information into sentences: he lost the ability to map thematic relations onto appropriate grammatical categories. Thus his broadly conceptual powers seemed normal, as did his understanding of bare lexical items; but his sentence-construction (and comprehension) abilities were severely impaired. This exemplifies intact thought (and subsentence understanding), but impaired sentence capabilities.¹⁴ In another case, certain patients described by Sirigu *et al.* (1998), who exhibited damage to Broca's area, could understand subsentential expressions such as 'pay', 'arrive at the news-stand', 'leave', and 'ask for the paper'. In particular, they could arrange these isolated phrases into 'scripts', choosing the appropriate temporal order. (In this example, the correct order for the script would be: 'arrive at the news-stand', 'ask for the paper', 'pay', 'leave'.) But despite understanding subsentences apparently as well as normal individuals, these Broca's aphasics were severely impaired when it came to ordering words *within a sentence*. (As the authors note, typical 'sentences' produced by these patients included 'The husband of less likeable my aunt is much than my cousin' and 'Gave a kiss to the boy the lady who it is'.) It is at least not obvious how this could be if people have to construct a complete sentence in order to understand an (apparently) subsentential utterance. (See also Varley 1998.)

Interestingly, this dissociation of thought and sentence abilities can go the other way as well. Thus DM, a patient described in Breedin and Saffran (1999), retained the ability to detect grammatical violations, and to recover

¹⁴ Of course, it is possible in principle that the person could *internally* construct sentences, even though she could neither produce them, nor distinguish grammatical ones from ungrammatical ones. But 'possible in principle' does not amount to plausible as an empirical hypothesis, unless some evidence is brought forward.

syntactic information generally—including especially the ability to assign thematic roles on the basis of grammatical clues. But DM had a severely degraded lexicon, no longer knowing the meaning of concrete words. Here we have impaired lexical (and thinking) skills, but relatively intact sentence capabilities.

None of these arguments is knock-down. But they do make the inner-speech theory of subsentential communication look implausible. And that is all we hoped to show here, since showing even this goes some way to answering (Q1) through (Q3):

- (Q1) How do hearers interpret subsentential speech of this sort?
- (Q2) How is perception-based information seamlessly combined with communication-based information?
- (Q3) In the light of the answer to (Q1), is it plausible that the phenomenon of subsentential speech is genuine, or is it more plausible that it is merely apparent?

The answers suggested by the foregoing arguments are these. First, regarding (Q1), hearers do not interpret subsentential speech by recovering a complete natural-language sentence: neither an ordinary sentence, nor an 'elliptical sentence'. Hence, regarding (Q2), it is not the case that perception-based information gets combined with communication-based information by always having the former encoded into an expression of natural language. For this is not what happens in subsentence cases of the sort described. Given this, the answer to (Q3) would seem to be that the phenomenon of subsentential speech is *not* merely apparent. For instance, our description of the 'From France' case at the beginning of the paper was accurate: Andrew produced a PP, whose meaning is a property, and Sylvia recovered that PP *and only it*. She then figured out the proposition asserted not semantically, by decoding a complete sentence, but pragmatically, by combining information got from perception with information got from her knowledge of language.

That said, to make the case for genuineness even more compelling, one needs a positive story in answer to (Q1), a way of fleshing out this rough sketch. We turn, therefore, to the two remaining initially plausible alternatives, namely (d) and (e).

4.3 *The Mentalese options, (d) and (e)*

Both remaining alternatives hold that the hearer tokens a Mentalese sentence, containing a Mentalese singular term referring to the displayed object. This is what it is to think about the object being discussed. In the Andrew example, the hearer tokens a term that stands for the cigarette, and she also tokens a Mentalese predicate which stands for that property had by all things from France. (She gets the latter by decoding the natural-language predicate produced by Andrew.) She puts these two Mentalese expressions together, and arrives at a Mentalese sentence, which encodes the aforementioned singular proposition.

In filling out this idea, the nature of Mentalese can be left open in numerous respects. For instance, whether it is mostly innate, whether it is universally shared by all humans *and* higher animals, etc. are not things we need to take a stand on. For our purposes, Mentalese is a language of thought, with compositional syntax and semantics—but no phonology. Hence, it is not a natural language. We shall also stipulate that the syntax of Mentalese is importantly different from that of spoken languages. What ‘importantly different’ amounts to is notoriously hard to say. It is fair to suppose that, unlike natural languages, Mentalese is disambiguated; fair also to suppose that all syntactic features in Mentalese have content effects. For instance, there will not be case markings that do not indicate grammatical function, or semantically vacuous prepositions, or verbal elements (like the copula, in some of its incarnations) whose contribution is merely grammatical. In a similar vein, Mentalese surely would not be ‘multi-stratal’: it would not, for instance, have analogues to Phonetic Form, D-Structure, S-Structure, and Logical Form. In contrast, natural languages are systematically ambiguous, they are riddled with grammatical features which are solely driven by surface syntax, and they are multi-stratal. (Why make this stipulation about being ‘importantly different’? The reason, at bottom, is that even assuming the truth of S-representationalism, it must remain a non-trivial claim that one thinks in Mentalese. If our Mentalese were, say, just like English but with no phonology, the claim would lose most of its interest.)

We assume a picture of the mind proposed and defended in Jerry Fodor’s 1983 book *Modularity of Mind*. We are not committed to the minutiae of

Fodor's views—nor, indeed, is he. But adopting this picture will help us express our speculations about subsentential communication fairly clearly. Given this picture, it seems impossible that the representation which is got by perception should remain, at every stage, in a format specific to the perceptual module in question. That is, considering the Andrew case again, the representation which stands for the cigarette in Andrew's hand cannot be one over which the visual system alone can operate, come what may. The obvious reason is that the representation must be capable of combining with another representation got by hearing and understanding a natural language phrase. Put in Fodor's terms, the two parts of the Mentalese sentence, at which the hearer eventually arrives, must be encoded in a single representational system. For, were that not the case, the representation of the cigarette could not be combined with the representation of the property λx . from-France(x)—because the latter representation comes, in the first instance, from a different module, the parser. But cross-modal understanding of this sort clearly does occur, since the hearer understands a proposition. Assuming this modular picture, then, it must be that *either* the representation of the cigarette is already in a format that can be conjoined with something got from the parser, *or* it is somehow 'translated' or 'incorporated' into that format prior to being combined with the representation that comes from the parser. (More on this in the final section.) This is one reason for thinking that the representation which encodes the whole proposition about the cigarette, to the effect that it is from France, is what Fodor calls a 'Central System' representation. There is, in addition, another reason for thinking this. As Fodor (1983) describes the Central System, it is the place where general-purpose inference occurs. But, as noted above, there are cases in which the propositions communicated in non-sentential speech serve as premisses in inferences. For instance, recall the imagined Francophobic debate, about whether anything really fabulous had been recently produced in France. In this instance, Sylvia draws inferences on the basis of the cigarette proposition, concluding that her claim, i.e. that nothing fabulous is being produced in France, is false. But then, if it is to serve as a premiss in an inference, the cigarette proposition must be in a format suitable for general-purpose inference. And that format, in Fodor's picture, is a Central System representation.

Let us now specifically consider the nature of the representation of the

cigarette within the Central System. The question we want to ask is: does the visual representation denote the cigarette *by description*, or is it rather an indexical whose content just is the displayed cigarette? That is, which of (d) or (e)—Descriptive Mentalese or Indexical Mentalese—in our options tree (Figure 8.1) is correct? We believe that there are broadly empirical considerations that favour the latter view. We turn to them now.

4.4 *Visual indexes; or, (e) over (d)*

This section is structured as follows. We begin by abstracting away from subsentential speech for the moment, considering how objects are 'visually grasped' in the general case. We first argue that if objects are 'visually grasped' via properties, then the only property which is at all viable is spatial location. We then appeal to recent experimental work in cognitive psychology to show that spatial location probably cannot do the job either. We conclude that, in general, objects are not visually attended to via their properties. The positive alternative, proposed by Zenon Pylyshyn in several recent writings, is that there are 'visual indexes', functionally reminiscent of natural-language demonstratives. Visually tracking an object is then a matter of tokening a visual index which refers to the object tracked. What this probably means for subsentential communication is this: it is surely plausible that, in answer to (Q1), visual indexes are among the things which are employed in understanding (certain cases of) less-than-sentential speech. That is, given how 'visual grasping' works in the general case, it seems reasonable to assume that, with respect to visual grasping in subsentential communication, it works in the same way. In which case (e), which posits indexicals, is the most promising option.

With respect to the more general question, abstracting away from subsentences, let us begin by considering whether non-locational properties would allow 'visual grasping'. The idea behind the proposal is this: what the person noticing an object does is token some kind of Mentalese definite description. The thing denoted is then the unique object which satisfies the (non-locational) properties specified in the definite description. There are some pretty obvious problems with this idea. First, there will not in general be just one thing which satisfies the description. So denotation is predicted to fail regularly, which it does not. Indeed, humans and other

animals have no trouble seeing numerous examples of the same kind of thing, distinguished only by their location. Thus we can very easily count the number of 'smiley' faces in the following line:



This despite the fact that the only difference between them is locational. (It is, of course, possible in principle that whenever we perceive 'exceedingly similar' objects, like the five smiley faces, there are very complicated and fine-grained—but non-locational—differences between them, which the visual system makes use of. If the descriptive material is that fine-grained, however, it is hard to see how we manage so easily to keep track of objects as they change their properties.¹⁵) These intuitive considerations aside, there are experimental results that point in the same direction, namely away from visual-system 'definite descriptions'—at least where the denotation is the unique thing exhibiting certain non-locational properties. For instance, take Figure 8.2. It is very easy to tell that the various shapes are in a linear order, and that there are six elements. What is more, experiments have shown that

¹⁵ There are other problems with the descriptivist approach. For one thing, the object of a singular judgement, whether perceptual or non-perceptual, figures essentially in the truth-conditions of that judgement. Suppose that what fixes the object of a thought is some descriptive condition that the object in question just happens to satisfy uniquely. Then, it is conceivable that the same thought could be true in worlds in which that particular object does not exist—as long as some other object satisfies the same condition in the world in question. In which case, the thought is not about the particular object in question (or about any particular object that happens to make its content true, for that matter). We can, of course, require that the relevant description be 'rigidified' by incorporating an 'actual' operator or some (rigid) designator of the actual world. Leaving aside the question of whether we normally think of objects in that way, the solution still does not help capture the *res* aspect of singular thoughts. To use Keith Donnellan's famous example, one can judge that Smith's actual murderer is insane and yet one's judgement fails to be about Jones, who actually murdered Smith. For it may be that one's judgement involves an attributive occurrence of the description in question (whether in some natural language or in Mentalese). Adding a causal relation to Jones (even a non-deviant one) will not help avoid that problem either, since the description may still be attributively used in that context. If it is, then the judgement would not be about Jones (strictly speaking); it would be about *whoever* murdered Smith, which in this case just happens to be Jones. For a comprehensive discussion of these and other related issues on singular thought, see Bach (1986); Burge (1977); and Lycan (1986).




FIG. 8.2. (from Pylyshyn 1999)

subjects can determine linearity and quantity of elements more quickly than they can judge whether any two elements are of the same type. With respect to the case at hand, you can more quickly judge that there are six elements in the above figure, and that they form a line, than you can judge whether only two items are the same. These facts suggest that subjects do not need to determine which objects satisfy which non-locational descriptions before they note *where* the objects are, and *how many* there are. Indeed, the subject's speed at recognizing linear order does not even depend upon how many objects there are, whereas, as Pylyshyn notes in a 1999 paper, later published in revised form in Branquinho (2001), the time it takes to judge whether the same display consists only of identical (or duplicated) objects increases linearly with the number of objects in the display. So, whereas recognizing which objects have which properties takes more time when there are more objects, recognizing what pattern the objects fall into does not take more time. But then doing the latter—seeing patterns, or quantity—cannot require doing the former, i.e. attending to descriptive properties had by each object. This strongly suggests that subjects do not have to identify the non-locational properties of each object when they count it, or see its place in a pattern. That is, they do not have to settle on a uniquely denoting description of the object. Since one cannot count objects, or see their place in a pattern, without grasping them, it seems equally unlikely that *grasping an object* requires describing it to oneself.

Having mentioned location, let us now consider whether the visual system, when it 'grasps an object', might be using 'definite descriptions' that employ *locational* properties. Appeals to intuition here are much less

helpful. So let us go straight to the experimental work. We quote at length from Pylyshyn:

In a typical experiment . . . observers are shown a screen containing anywhere from 8 to 24 simple identical objects (points, plus signs, circles, figure-eight shapes). A subset of these objects is briefly rendered distinct (usually by flashing them on and off a few times). Then all the objects (which are visually identical) move about in the display in unpredictable ways. The subject's task is to keep track of this subset of objects (called 'targets'). At some later time in the experiment (say 10 seconds into the tracking trial) one of the objects is again flashed on and off. The observer must then indicate whether or not the flashed (probe) figure was one of the targets. (In other studies the subject has to indicate *all* the targets using a mouse or other pointing device.) A large number of experiments, beginning with the studies described in Pylyshyn & Storm (1988), have shown clearly that observers can indeed track up to 5 independently moving identical simple visual objects. (Pylyshyn 1999: 13–14)

This finding is important because computer simulations of object-tracking which employed definite descriptions that denoted via locational descriptions did not come even close to human performance. The locational-/description-based simulations made many more mistakes, regularly losing track of the objects.

To get an intuitive handle on the proposal to be rejected, consider a little thought-experiment. Imagine that each object in our visual system is described via an ordered triple, using Cartesian co-ordinates. That is, each is described as 'object at co-ordinate x, y, z '. In Pylyshyn's experiments, we start off with about five objects tracked, plus 15 in the background. So, in the imagined visual system, that would be 20 locational triples. With each motion of each object, all the co-ordinates would have to be updated. And such a system would need a tracking function that would, at least five times over, have to note something like: ' $O_{\langle 28, 132, 87 \rangle}$ at $t_0 = O_{\langle 29, 133, 87 \rangle}$ at t_1 '. The complexity of such an imaginary tracking system is daunting. Now, obviously one need not suppose that the locational descriptions, if any, which humans actually use are precisely like these imagined ones. But it gives one the idea of what a locational-description view would say.

Now back to reality: Pylyshyn's simulations, using his sophisticated locational descriptions, produced lots of errors, far more than real humans, who actually do quite well at tracking objects in the visual field. Assuming

that his simulations were otherwise accurate, it seems safe to conclude that we humans—who succeed so well in this task—do not in general attend to objects via a locational description. (Surprisingly too, subjects were able to keep track of up to five objects as they changed location *and their other properties*.) Pylyshyn and others have concluded that descriptions of location cannot be what humans actually use as the basis for visual tracking. But, given that, as argued above, locational descriptions were the best candidate for an identifying property, we conclude that objects are attended to without using any description of them, locational or otherwise.

To be sure, Pylyshyn's experiments do not fully justify the claim that visual indexes function as referential devices that refer non-descriptively to perceived *objects*. For instance, the results described above are also compatible with the claim that, in vision, we use *spatial locative demonstratives* ('there' rather than 'that') in the visual system. That is, for all we have said here, the visual system might refer pre-conceptually to the *locations* of objects, then 'describing' the object via such demonstratives. On this view, visual indexes are referential mechanisms that directly pick out locations in the subject's visual field, without either referring to or describing the objects that occupy those locations. Visual indexes non-descriptively encode information about the spatial locations of moving or stationary objects. As such, they function as indexical representations of spatial locations rather than as representations of spatially located objects. (In terms of the thought-experiment introduced above, the object would then be not $O_{(x,y,z)}$ but O_{there} . It is clear how this would simplify the computational burden.) We therefore concede that Pylyshyn's experiments do not conclusively rule out the view that visual indexes function as spatial locatives. That is an empirical matter on which, given our immediate purposes, we need not take a particular stand.¹⁶ But, as before, Pylyshyn's own view emerges as initially the most plausible. (It is worth recalling that if *either* option (d) or (e) ends up being true, then subsentential speech is surely genuine. So the outstanding issues at this point, when applied to subsentence cases, have to do with (Q1) and (Q2), not (Q3).)

¹⁶ Pylyshyn's general position is not undermined if it should turn out that visual indexes function as demonstrative spatial locatives rather than as demonstratives that pick out objects. For his experiments are intended to show only that we do not perceptually track things solely by describing them.

Following Pylyshyn, then, we conclude that there are 'visual indexes'. These are representations within the visual system which refer to, without describing, objects in the environment.¹⁷ It is worth stressing that visual indexes are representations whose contents are determined only relative to their context of occurrence. (Cf. Bach 1986; 1987; and Récanati 1993.) Two individuals may token type-identical visual indexes, in qualitatively indistinguishable perceptual environments, have qualitatively indistinguishable percepts, and yet have different perceptions because the contextually determined objects of the tokened indexes are distinct. The object of a (tokened) visual index is not thought of as the unique object of a certain sort, i.e. under a definite description (whether in Mentalese or in a natural language), but is rather thought of in some contextually sensitive way. To put the point another way, something is an object of perception not in virtue of its uniquely satisfying some descriptive condition imposed by the perceiver's representation of her environment; but rather, it is an object of perception in virtue of its standing in the appropriate causal-contextual relation to the perceiver, a relation that the person's perceptual mechanisms do not represent in thought.

We propose to treat visual indexes as representations in the visual system rather than as outputs of the visual system. Their contents, which are the perceived objects, get into the Central System by further processing. Although we do not know how the process actually works, we shall argue in the next section that the representational contents of visual indexes somehow combine with material got from the parser in the understanding of non-sentential speech—at least in the sort of case considered here. If that is right, then, to account for subsentence cases, there must be representations in the Central System that somehow link up with visual indexes in some indirect way. We shall turn now to that topic, taking for granted from here on that, because of how visually grasping seems to work in the general

¹⁷ An early version of the idea that subjects mentally employ visual representations as singular referential devices can also be found in Chastain (1975). However, Chastain thinks that it is the conscious *visual experience* involved in a subject's seeing an object—and in virtue of which it is (in that context) a seeing of *that* object—that actually functions as a singular referring mechanism. Neither Pylyshyn nor we are committed to that view since all three of us hold that visual indexes function at a level of processing inaccessible to conscious experience or direct awareness.

case, option (e) from Figure 8.1 above is the most viable suboption for how hearers understand subsentential speech.

5 From Visual Indexicals to Indexicals in Mentalese

A quick review, before we proceed. We began by introducing a single example of subsentential communication: Andrew uttered the prepositional phrase 'From France', and thereby asserted, of a visually salient cigarette, that it was from France. We then noted a host of philosophical doctrines which, at least at first glance, appear to be in tension with the existence of subsentential speech acts. These included theses about: the domain of logical form, the contextual determinants of what is said, the sufficiency of knowledge of language in interpreting literal speech, the proper analysis of assertion, the location of the semantics/pragmatics boundary, the context principle and its implications, etc. Given the importance of these doctrines for philosophy of language and epistemology, two issues become pressing. First, is subsentential communication a genuine phenomenon? Second, if it is genuine, are the philosophical doctrines described genuinely in tension with it? We have skirted the latter question. But we have been trying to urge a positive answer to the former question, largely by considering five options for treating subsentential speech. One result of that investigation was that the options which treat the phenomenon as *merely* apparent face real problems. (These were options (b) and (c). Option (a), it will be recalled, was put aside for future research.) A striking upshot of this, with respect to at least one larger philosophical doctrine, was that thought is not inner speech. Or anyway, occurrent thought does not require that a natural-language vehicle for that (whole) thought 'pass through the mind'. This in itself is not so surprising. It has long been suggested that language learning, image manipulation, task memory, etc. do not involve inner speech. But if we are right, not even reasoning, nor argumentation, *nor even the understanding of speech* can be exhaustively explained by inner speech of this sort. The reason is, the thought which the hearer recovers and uses as a premiss, in subsentential communication, goes well beyond the meaning of the words which pass through her mind. Closer to the truth, at least if the S-representational picture is correct, is that the recovery of the thought involves the processing of sentential mental representations of a special sort, Mentalese expressions.

Having arrived at that result, we began to delve more deeply into the nature of the representations processed during subsentential comprehension. Part of the point of this, beyond its intrinsic interest in allowing us to address (Q1) and (Q2), is to support the genuineness of subsentential understanding—by telling a positive story, consistent with what is known, that shows it to be such. To develop this story, we abstracted from subsentential speech for a moment, and explained and defended the notion of a ‘visual index’, a device which allows the visual system to track objects indexically. The plausibility of this idea in the general case of ‘visually grasping’ supported option (*e*), an option according to which the phenomenon is genuine. We now pursue (*e*) even further.

The main question of this section is this: what precise bearing does our discussion of visual indexes have on the issue of non-sentential communication? In particular, how might visual indexes get employed in subsentential understanding? To answer that question, we shall borrow some helpful ideas from Joseph Levine’s account of indexical thought within Mentalese (cf. Levine 1988). His account operates on the assumption that we think in Mentalese, which is also our working hypothesis. To see how his account works, let us first consider a case in which a speaker explicitly uses a demonstrative to refer to some perceived object. Suppose you see a dog a few yards away, and you say, ‘That is a cocker spaniel’, as you point to the dog. You thereby expressed a singular, demonstrative, thought about the dog. Levine’s question is this: What is the cognitive role of the linguistic demonstrative in the expression of an indexical *thought* about the dog? In developing his answer to that question, Levine presents the following picture:

our cognitive architecture is modular at least to the extent that we have relatively autonomous perceptual systems which take impinging stimuli as inputs and deliver perceptual representations—percepts—as outputs. Percepts then serve as inputs to the more general ‘CPU’, the system within which belief-fixation takes place. A mental demonstration is a particular kind of causal/computational relation between an expression of mentalese and a percept. The mentalese expression standing in this relation to a percept is a pointer. We can think of a pointer as an ‘internal finger’ pointing at the percept, except that the ultimate object of the demonstration is not the percept ‘ostended’ itself, but the external object it represents . . . The main point is that some computationally significant causal relation

obtains between the mental representation corresponding to the use of a demonstrative and a perceptual representation output from the perceptual module such that the former inherits its referent from the latter. (Levine 1988: 232)

Levine continues:

A pointer, then, is a mentalese demonstrative that underlies my use of demonstratives in natural language. The canonical form of a pointer can be represented as follows: $[\rightarrow(\alpha)]$, where α designates a percept. A pointer is really a directly referential term, its referent being the object represented by the percept α .

Thus, with respect to the present example, the indexical thought you expressed when you reported what you saw ('That is a cocker spaniel') involves the generation of a perceptual representation of the dog. We shall use α to designate your percept of the dog rather than the dog itself. To facilitate discussion, we shall depart from Levine's notation by representing Mentalese indexicals—whether demonstrative or otherwise—as English indexicals in bold capitals. Thus, in the dog example, a token of $IT_{[\alpha]}$ (in your language of thought) refers to the dog you saw. Being a perceptual case, it does so in virtue of demonstrating a percept, namely α , that itself represents the dog. The demonstrated percept, in turn, represents the dog in virtue of certain causal facts about the perceptual situation, including facts about the visual indexes that play a causal role in the tracking of the perceived dog. The indexical thought that you mentally tokened in your Belief Box, and which you expressed by means of your report, was: **COCKER SPANIEL**($IT_{[\alpha]}$).¹⁸

Levine's account of the cognitive role of demonstratives in natural language is quite ingenious. We would like to make two points about it. First, it is entirely neutral on the empirical question of the nature, structure, and representational properties of percepts—as it should be. In particular, his account of indexical thought is compatible with the empirical hypothesis, supported above, that we employ visual indexes when we visually perceive objects. Nothing in the account rules out the possibility that, in some perceptual contexts, and at some level of processing, vision employs visual

¹⁸ We wish to remind the reader of our convention of using capitalized English expressions in bold to represent Mentalese expressions. We do not mean to suggest that Mentalese is a natural language. On the contrary, we have supplied an indirect reason in this paper to think that it is not.

indexes. Thus, the 'visual index' account can be made to fit nicely with the view we are outlining in this section.

Second, even though Levine's account is an explanation of the cognitive role of public-language demonstratives in the expression of demonstrative thought, we see no reason why it cannot also be used to explain how hearers can cognitively grasp singular thoughts that speakers often convey by way of non-sentential speech. The difference is that, in the latter case, the speaker does not actually produce a natural-language demonstrative. Nor is her utterance 'elliptical' for a natural-language sentence that contains a demonstrative (whether a pure demonstrative, a complex demonstrative, or some other indexical expression). In the dog example, you could have drawn attention to what you saw by assertively uttering the bare phrase 'Cocker spaniel' as you pointed to the dog, rather than by uttering 'That is a cocker spaniel'. And yet, your friend will have interpreted you as having said of the demonstrated dog that it is a cocker spaniel.

To explain how both things can be true, let us apply Levine's story, thereby illustrating our answer to (Q1). In interpreting your utterance, your friend employs a Mentalese indexical that designates (in that perceptual context) the dog you pointed to. However, it refers to the dog by virtue of demonstrating, via some causal and computational route, a (tokened) perceptual representation your friend formed (presumably a visual index) of the dog. This, we may suppose, is how visual indexes manage to get 'translated' into the Central System: what really happens is that a (tokened) Mentalese indexical demonstrates a percept, whose represented object is indexically fixed by a visual index, and takes over as its referent the referent of that index. Eventually, at some level of Central System processing, your friend stores a Mentalese sentence of the form:

SAID(S)(COCKER SPANIEL(IT_[α]))

in his Belief Box, where **S** is a Mentalese expression that refers to the speaker; **α** is a percept, i.e. a representation within the visual system, specifically a visual index that (in this context) designates the dog; and **IT_[α]** is a Mentalese indexical that (in this context) refers to the dog indirectly, by virtue of demonstrating the visual index **α**.

If true, this explains how a hearer can grasp, without tokening a natural-language sentence, a singular thought conveyed by a speaker who utters a

subsentential expression. And, we stress, it does not follow from this story that the assertive utterance of the phrase 'Cocker spaniel' is semantically or syntactically equivalent to (or elliptical for) an utterance of the sentence 'That is a cocker spaniel', even if the same speaker reference/demonstratum is involved. Hence, in answer to (Q₃), the genuineness of the phenomenon is maintained. At most, all that follows (if our extension of Levine's account to non-sentential assertions is correct) is that, in certain communicative contexts, a bare phrase and a natural-language sentence can have the same underlying cognitive role within the Central System, in speech perception and interpretation. (This is not surprising if, as seems inevitable, natural-language sentences themselves must get translated into Mentalese during interpretation.) Similar remarks apply, of course, in the case of assertive uses of prepositional phrases. Consider, once again, our example about Andrew and his cigarette. By hypothesis, the hearer interpreted Andrew's non-sentential utterance of 'From France' as meaning that *that* cigarette (the one that he demonstrated) is from France. Our proposal is that, in understanding Andrew's utterance, the hearer used a Mentalese indexical, which we shall represent as $IT_{[\beta]}$ that refers (in that context) to the visual index β , and thereby refers, via deferred demonstration, to the cigarette. Computationally, the tokening of the Mentalese indexical demonstrates, in Mentalese, the (tokened) percept of the cigarette; semantically, it refers to the cigarette.¹⁹ It is also combined with the representation that the language module emits from the processing of Andrew's verbal input, namely, the Mentalese predicate **FROM-FRANCE**. The result is (something like) the Mentalese sentence **FROM-FRANCE**($IT_{[\beta]}$). That sentence is, in turn, encoded as part of a more complex Mentalese sentence,

SAID(ANDREW)(FROM-FRANCE($IT_{[\beta]}$)),

which Sylvia comes to harbour in her Belief Box. Semantically, the latter

¹⁹ Relative to the context in question, the tokened Mentalese demonstrative, $IT_{[\beta]}$, does not refer to the percept β . It refers instead to Andrew's cigarette. But it does not refer to it by describing the cigarette in terms of certain properties that the cigarette uniquely possesses. Rather, it refers to the cigarette by its demonstration of a (tokened) percept that represents the cigarette. We are suggesting, then, that Mentalese indexicals, at least in these perceptual cases, function as deferred demonstratives in the language of thought. For a defence of the view that deferred demonstratives in natural language are directly referential singular terms, see Borg (2001).

sentence is true just in case, relative to the context in which he displayed the cigarette and uttered 'From France', Andrew said of the cigarette that it is from France. If Andrew said that, then it is reasonable to conclude that he asserted a singular proposition, one that (necessarily) is true if and only if that particular cigarette is from France. It is also reasonable to conclude that that is what Sylvia interpreted him as saying—even though she did not recognize any available natural-language sentence in interpreting his utterance.

6 Conclusions and Further Research

We can best summarize our conclusions by recalling the three guiding questions of the paper:

- (Q₁) How do hearers interpret subsentential speech of this sort?
- (Q₂) How is perception-based information seamlessly combined with communication-based information?
- (Q₃) In the light of the answer to (Q₁), is it plausible that the phenomenon of subsentential speech is genuine, or is it more plausible that it is merely apparent?

Recalling that we were only discussing the use of predicates to talk about objects that are visually grasped, our answer to (Q₁) goes like this. The hearer forms a visual index, which is locked onto the displayed object. That index, in turn, is linked up with a Mentalese indexical, where the said linking is achieved by demonstration of the visual index. As a result of this two-step process, the Mentalese indexical derivatively refers to the displayed object. The Mentalese indexical then gets combined with a Mentalese predicate, where the latter comes from the language module, having been decoded from the natural-language predicate actually spoken. It is, we conjecture, the concatenation of such a parser-derived Mentalese predicate with a Mentalese indexical, whose reference is contextually fixed by a visual index, that constitutes understanding less-than-sentential speech, at least in cases of visually grasping objects, and combining them with a predicate-meaning got from the parser. (Or anyway, it is something like that.) What this suggests about (Q₂) is that at least in these cases the information is integrated by having both kinds of information converted into Mentalese—the one from

visual processing, the other from linguistic decoding. What is *not* going on is the translation of all the information into natural language. Finally, what all of this suggests as an answer to (Q₃) is that the phenomenon is indeed genuine—it *appears* that neither the speaker nor the hearer processes a natural-language sentence, because neither of them *does* process a natural-language sentence.

Of course there remain two questions that, as we warned above, we have not addressed here—even putting aside the fact that the issue of option (a) in Figure 8.1, the non-S-representationalist option, remains unresolved. The first of these is the question of whether the apparent tensions between sub-sentential speech and the various philosophical doctrines noted in Section 2, are real tensions. The second is equally important. We end with it.

What, one surely wants to know, is to be said about cases of non-sentential speech that are not as simple as the one described here? We have been discussing only perceptual cases—indeed, only perceptual cases involving vision. But it is quite clear that there must be more to the story than that. In addition to their *visual demonstrative* roles, Mentalese indexicals can serve other functions. For a start, in the same sense that one can ‘visually demonstrate’, one can ‘demonstrate’ an object by hearing it, or smelling it, or otherwise sensing it. And these cases too will give rise to subsentence use. Altering the Andrew case somewhat, for example, one might notice the quite distinctive scent of some cigarette smoke, and say: ‘From France’. Here the cigarette being described, i.e. the one which is the source of the smell, is not linked up to the Mentalese demonstrative by a *visual* index. For the cigarette has not (yet) been seen. Or again, hearing a very loud noise, recognizable as a sonic boom, Rohit could say to Dom: ‘An American jet fighter’, thereby describing not the boom, but its source.

Moreover—non-visual perceptual cases aside—one can talk about an object which is not *presently* perceived at all, visually or otherwise. Thus, suppose that several days after the just-mentioned sonic boom, Dom reads that the Concorde has begun flying over their town, throttling up to Mach 1 precisely overhead. He could then say to Rohit, at their next meeting: ‘That boom the other day? Not a jet fighter. The Concorde’. For such cases, one would need Mentalese indexicals whose content is not determined demonstratively. They would be multi-purpose, usable for any contextually salient thing however it is grasped. Demonstration of a percept would apply only

in the specific case of presently perceived objects. For instance, Mentalese indexicals would need to serve as temporary file-names for contextually determined referents, and be maintained even when the referents are not perceptually available. Information about the object could then be appended to the symbol as a way of keeping a track history of the object. (Returning to the first example, the hearer's perception of Andrew's cigarette resulted in her acquiring and storing information about it in short-term memory, so that she can later come to wonder where Andrew bought the cigarette, even when no longer perceiving it. The reason one can do this, we conjecture, is that there is now a 'file' which stands for that object, available for later use. This would also help explain, by the way, why one is able to go from hearing Andrew's assertive utterance of 'From France', and seeing the demonstrated cigarette, to reporting things like, 'Andrew said that the cigarette he was holding in his right hand is from France', 'Andrew said that that cigarette is from France', etc., even though Andrew did not use any of these singular terms.) Indeed, Mentalese indexicals can surely be used even if the object discussed was not perceptually grasped in the first place: the object discussed could be a conjectured object. Here is an example. Two physicists could be discussing the strange behaviour of a newly postulated particle, the Fridgeon. Returning the next day, one scientist could say to the other: 'Positron-emitting'. She would thereby have asserted, about the *postulated entity*, not perceived by any of the senses ever, that it is positron-emitting.

One last complication. As has just emerged, it is emphatically not the case that subsentences can be used only to talk about visually perceived objects: one can also speak subsententially about objects perceived non-visually (the smelt cigarette, the heard jet), and even about things never perceived. That the thing talked about happened to be visually perceived is, therefore, a quite accidental feature of the original Andrew case. But, surprisingly, that it is a thing—i.e. an object—being talked about is *also* an accidental feature of that specific example. (Though it is not accidental that we selected such an example: the aim of the paper was to take the very first steps towards understanding how subsentential speech is understood. Hence we deliberately picked a very simple case.) One can perform subsentential speech acts in which 'what is salient'—seen, heard, remembered, imagined, whatever—is a property, or even a generalized quantifier. Thus, one can point at graffiti on a wall and say 'John', meaning thereby that John

did the *graffiti*. (Here, a property is supplied by the environment, and an object is supplied by decoding.) Or, one could say to a person who is searching everywhere in the house, clearly in need of a cigarette, 'In the fridge', thereby asserting that *there are cigarettes* in the fridge.²⁰ (Here a generalized quantifier is contextually supplied, with decoding providing a missing property. For more on quantificational cases of subsentential speech, see Barber, forthcoming; Botterell, forthcoming; Stainton 1998a; 1998b.) Clearly, then, there is *much* more to be said about these matters. But that will have to wait for another day.

REFERENCES

- BACH, K. (1986), 'Thought and Object: *De Re* Representations and Relations', in Brand Harnish (1986), 187–218.
- (1987), *Thought and Reference* (Oxford: Oxford University Press).
- (1994a), 'Semantic Slack: What is Said and More', in Tsohatzidis (1994), 267–91.
- (1994b), 'Conversational Implicature', *Mind and Language*, 9: 124–62.
- BARBER, A. (forthcoming), 'Co-extensive Theories and Unembedded Definite Descriptions', in Elugardo and Stainton (forthcoming, b).

²⁰ Further complicating things, Ernie Lepore drew our attention to very interesting modal cases. Thus Andrew could have said not 'From France', but 'Maybe from France'. It is not obvious what the process of interpretation of such utterances would be, given our answer to (Q1). It might seem, for instance, that the modal 'maybe' must apply to a complete sentential structure—something which, on our view, is only available *after* the complete Mentalese sentence, **FROM-FRANCE**(IT_(a)), is constructed. But then the language module cannot be outputting a unitary meaning for the element uttered, since there isn't the right sort of thing for 'maybe' to modify: it modifies complete sentential expressions, not prepositional phrases. We remain unsure how exactly to treat this sort of example. As a first pass, however, we would note that 'maybe' and other modals can actually modify subsentential structures within sentences. For example, suppose the police suspected a six-foot-two criminal in a robbery. A witness describing the robber could disabuse them of their suspicion by saying: 'The bandit was about five foot eight, maybe six foot, but definitely not six foot two'. Here 'maybe' does not modify the complete proposition. Nor does 'definitely not' modify that proposition. Rather, the thing modified is subpropositional. Our hope, then, is that whatever story is told about modal modifiers in these less-than-propositional cases will equally apply to modals used in subsentential speech.

- BARTON, E. (1990), *Nonsentential Constituents* (Philadelphia: John Benjamins).
- BERMAN, S., and HESTVIK, A., eds. (1992), *Proceedings of the Stuttgart Ellipsis Workshop* (Sprachtheoretische Grundlagen für die Computerlinguistik, Bericht 29; Stuttgart: University of Stuttgart).
- BORG, E. (2001), 'Deferred Demonstratives', in Campbell, O'Rourke, and Shier (2001), 214–30.
- BOTTERELL, A. (forthcoming), 'Acquaintance and Isolation', in Elugardo and Stainton (forthcoming, b).
- BRAND, M., and HARNISH, R. M., eds. (1986), *The Representation of Knowledge and Belief* (Tucson: University of Arizona Press).
- BRANQUINHO, J., ed. (2001), *The Foundations of Cognitive Science* (Oxford: Clarendon Press).
- BREDIN, S. D., and SAFFRAN, E. M. (1999), 'Sentence Processing in the Face of Semantic Loss', *Journal of Experimental Psychology: General* 128: 547–62.
- BURGE, TYLER (1977), 'Belief *De Re*', *Journal of Philosophy*, 74: 338–62.
- CAMPBELL, J. K., O'ROURKE, M., and SHIER, D., eds., (2001), *Meaning and Truth: Investigations in Philosophical Semantics* (New York: Seven Bridges Press).
- CARRUTHERS, P. (1996), *Language, Thought, and Consciousness* (Cambridge: Cambridge University Press).
- and BOUCHER, J., eds. (1998), *Language and Thought: Interdisciplinary Themes* (Cambridge: Cambridge University Press).
- CARSTON, R. (1988), 'Implicature, Explicature, and Truth-Theoretic Semantics', in Kempson (1988), 151–81.
- (2002), *Thoughts and Utterances* (Oxford: Blackwell).
- CHAO, WYNN (1988), *On Ellipsis* (New York: Garland).
- CHASTAIN, C. (1975), 'Reference and Context', in Gunderson (1975), 194–269.
- CHATTERJEE, A., MAHER, L. M., GONZALEZ ROTH, L. J., and HEILMAN, K. M. (1995), 'Asyntactic Thematic Role Assignment: The Use of a Temporal-Spatial Strategy', *Brain and Language*, 49: 125–39.
- CLAPP, L. (2001), 'What Unarticulated Constituents Could Not Be', in Campbell, O'Rourke, and Shier (2001), 231–56.
- (forthcoming), 'On the Interpretation and Performance of Non-Sentential Assertions', in Elugardo and Stainton (forthcoming, b).
- DAVIS, S. (forthcoming), 'Quinean Interpretation and Anti-Vernacularism', in Elugardo and Stainton (forthcoming, b).
- DEVITT, M., and STERELNY, K. (1999), *Language and Reality*, 2nd edn. (Cambridge, Mass.: Bradford Books/MIT Press).

- DUMMETT, M. (1973), *Frege: Philosophy of Language* (Cambridge, Mass.: Harvard University Press).
- ELUGARDO, R., and STAINTON, R. J. (2001), 'Logical Form and the Vernacular', *Mind and Language*, 16: 393-424.
- (forthcoming, a), 'Shorthand, Syntactic Ellipsis, and the Pragmatic Determinants of What is Said', *Mind and Language*.
- eds. (forthcoming, b), *Ellipsis and Non-Sentential Speech* (Dordrecht: Kluwer).
- FODOR, J. A. (1983), *Modularity of Mind* (Cambridge, Mass.: MIT Press).
- GAUKER, C. (1997), *Thinking Out Loud: An Essay on the Relation between Thought and Language* (Princeton: Princeton University Press).
- GUNDERSON, K., ed. (1975), *Language, Mind, and Knowledge* (Minnesota Studies in the Philosophy of Science, 7; Minneapolis: University of Minnesota Press).
- HANKAMER, J., and SAG, I. (1976), 'Deep and Surface Anaphora', *Linguistic Inquiry*, 7: 391-426.
- KACHRU, B. B., LEES, R. B., MALKIEL, Y., PIETRANGELI, A., and SAPORTA, S., eds. (1973), *Issues in Linguistics: Papers in Honor of Henry and Renée Kahane* (Urbana, Ill.: University of Illinois Press).
- KEMPSON, R. M., ed. (1988), *Mental Representations* (Cambridge: Cambridge University Press).
- KENYON, T. (1999), 'Non-Sentential Assertions and the Dependence Thesis of Word Meaning', *Mind and Language*, 14: 424-40.
- (forthcoming), 'Non-Sentences, Implicature, and Success in Communication', in Elugardo and Stainton (forthcoming, b).
- LAPPIN, S., and BENMAMOUN, E. (1999), *Fragments: Studies in Ellipsis and Gapping* (Oxford: Oxford University Press).
- LEVINE, J. (1988), 'Demonstrating in Mentalese', *Pacific Philosophical Quarterly*, 69: 222-40.
- LUDLOW, P. (1999), *Semantics, Tense and Time*, Cambridge, Mass.: MIT Press.
- (forthcoming), 'A Note on Alleged Cases of Nonsentential Assertion', in Elugardo and Stainton (forthcoming, b).
- LYCAN, WILLIAM G. (1986), 'Thoughts about Things', in Brand and Harnish (1986), 160-86.
- (1988), *Judgment and Justification* (Cambridge: Cambridge University Press).
- (1990), 'Mental Content in Linguistic Form', *Philosophical Studies*, 58: 147-54.
- MORGAN, J. (1973), 'Sentence Fragments and the Notion "Sentence"', in Kachru et al. (1973), 719-51.

- (1989), 'Sentence Fragments Revisited', *Chicago Linguistics Society: Papers from the Parasession on Language in Context*, 25: 228–41.
- PYLYSHYN, Z. (1999), 'Connecting Vision with the World: Evidence for a Pre-conceptual Mechanism that Individuates and Keeps Track of Objects' (draft manuscript, Rutgers University).
- and STORM, R. W. (1988), 'Tracking of Multiple Independent Targets', *Spatial Vision*, 3: 1–19.
- RÉCANATI, F. (1989), 'The Pragmatics of What is Said', *Mind and Language*, 4: 294–328.
- (1993), *Direct Reference: From Language to Thought* (Oxford: Blackwell).
- SAG, I. (1976), 'Deletion and Logical Form' (diss. Ph.D., Massachusetts Institute of Technology).
- SELLARS, W. (1963), *Science, Perception, and Reality* (London: Routledge and Kegan Paul).
- SIRIGU, A. (1998), 'Distinct Frontal Regions for Processing Sentence Syntax and Story Grammar', *Cortex*, 34: 771–8.
- SPEERBER, D., and WILSON, D. (1986), *Relevance: Communication and Cognition* (Cambridge, Mass.: Harvard University Press).
- STAINTON, R. J. (1995), 'Non-Sentential Assertions and Semantic Ellipsis', *Linguistics and Philosophy*, 18: 281–96.
- (1997a), 'Utterance Meaning and Syntactic Ellipsis', *Pragmatics and Cognition*, 5: 49–76.
- (1997b), 'What Assertion is Not', *Philosophical Studies*, 85: 57–73.
- (1998a), 'Quantifier Phrases, Meaningfulness "in Isolation", and Ellipsis', *Linguistics and Philosophy*, 21: 311–40.
- (1998b), 'Unembedded Definite Descriptions and Relevance'. *Revista ali-cantina de estudios ingleses* (special issue on Relevance Theory), 11: 231–9.
- (2000), 'The Meaning of "Sentences"', *Noûs*, 34: 441–54.
- (forthcoming), 'In Defense of Non-Sentential Assertion', in Szabo (forthcoming).
- STALNAKER, R. (1990), 'Mental Content and Linguistic Form', *Philosophical Studies*, 58: 129–46.
- STANLEY, J. (2000), 'Context and Logical Form', *Linguistics and Philosophy*, 23: 391–433.
- SZABO, Z., ed. (forthcoming), *Semantics vs. Pragmatics* (Oxford: Oxford University Press).
- TRAVIS, C. (1985), 'On What is Strictly Speaking True', *Canadian Journal of Philosophy*, 15: 187–229.

TSOHATZIDIS, S. L., ed. (1994), *Foundations of Speech Act Theory* (London: Routledge).

VARLEY, R. (1998), 'Aphasic Language, Aphasic Thought: An Investigation of Propositional Thinking in A-Propositional Aphasic', in Carruthers and Boucher (1998), 128–45.

WILLIAMS, E. (1977), 'Discourse and Logical Form', *Linguistic Inquiry*, 8: 101–39.