2 From outer to inner space: linguistic categories and non-linguistic thinking

Stephen C. Levinson

Max Planck Institute for Psycholinguistics

Introduction

What is the relation between the medium in which we think and the medium in which we talk? Opinions have varied with intellectual character and climate: some, from Wilhelm von Humboldt to Benjamin Whorf, from Lev Vygotsky to Daniel Dennett, have imagined language to be the great facilitator; others, like Bertrand Russell or Gilbert Ryle (or for that matter Alfred Korzybski), have conceived of language as a veil of confusion that hides the cognitive essentials. On top of this, some have assumed that languages will all have the same underlying semantics; others have assumed they will differ deeply at a conceptual level. These opinions have all the hallmarks of speculation without empirical constraint, and one could be forgiven for thinking that the subject should be left, like other phantoms, to the philosophers. But I believe that we can at least hedge in the fantasies by working in two directions. The first is to apply first principles to a few crucial facts about language. The second is to treat the problem as an empirical one: very well, let us roll up our sleeves and investigate whether a difference in linguistic conceptualization is or is not correlated with a difference in pattern of thinking. In this chapter, I sketch these two tacks in the last section, hoping to suggest that real progress could be made, while in the first section I try to sweep the ground clear of obvious misconceptions.

1 Never the twain shall meet: why semantic representations are distinct from conceptual representations but after all necessarily related

1.1 Half a dozen reasons why

Do we think the way we speak? Scholars are wont to transform this simple question into something much more complex and perhaps not much clearer. Nevertheless, they divide sharply into lumpers (let us call them

13

A-theorists), who see no reason for a distinction between semantics and underlying conceptual representations, and splitters (let us call them B-theorists), who insist on a distinction. Let me introduce you to some of each.

Some A-theorists

"Cognitive grammar therefore equates meaning with conceptualization . . . " (Langacker 1987: 5).

"The terms *semantic structure* and *conceptual structure* denote the same level of representation" (Jackendoff 1983: 95).

"Perhaps the reason that semantic representations have proved to be so elusive is that, after all, there aren't any" (Fodor et al. 1975: 530).

Some B-theorists

"Linguistically encoded semantic representations are abstract mental structures which must be inferentially enriched before they can be taken to represent anything of interest" (Sperber & Wilson 1976: 174).

"Understanding an utterance is knowing a proposition that in the context the hearer CONSTRUCTS from the meaning of a univocal, sense-general sentence. The general meaning of the sentence is made specific in the hearer's understanding of the utterance" (Atlas 1989: 146).

"Meaning underdetermines interpretation" (Barwise & Perry 1983: 37).

"The way in which experience is conceptualized is no simple isomorphism of the way in which it is verbally expressed . . . Thus SF (Semantic Form) tends to be conflated with CS (Conceptual Structure), very much like phonology tended to be conflated with articulatory phonetics. There are, however, strong reasons to distinguish SF and CS" (Bierwisch & Schreuder 1991: 24; 30–31).

"A cognitive psychologist might take these puzzles as arguments for the view that what is retained in memory is not retained either in linguistic form . . . or in propositional form . . . Perhaps what is retained is a mental representation of the episode itself – a representation that is non-linguistic, yet adequate to support linguistic representations, with or without deictic terms at a later date" (Miller 1982: 66–67).

"The idea that thought is the same thing as language is an example of what can be called a conventional absurdity . . . there has to be a 'what we meant to say' that is different from what we said" (Pinker 1994: 57).

Yet one has the suspicion that some of these theorists have often implicitly assumed answers to our question without drawing out the consequences. Let us try to do so. But there are layers beneath layers of assumptions here, and we cannot do the archaeology of the whole stratified heap. Let us then adopt without too many scruples the parlance of the representational theory of mind, the viewpoint of 'classical' cognitive science, merely noting that 'representation' tends to come biased towards the language-like rather than the imagistic or pictorial, and that there's a radical deflation of all this in connectionist circles (see e.g. Sterelny 1990 for a dispassionate overview). But the notion of 'representation' raises a caveat: different scholars have deeply different views about what kind of a thing a representation of a thought or a meaning might be. The exegesis of such different (alas, often incoherent) views is none of our business here – we are concerned simply with boundary conditions, specifically with whether the representation of a thought and a meaning could in principle be the same sort of thing, whatever it is. Thus for these purposes I shall be happy to group scholars who under a broader light would be the strangest of bedfellows.

So, on these minimal assumptions, our simple question then gets transformed from 'Do we think the way we speak?' to 'Are the representations that underlie linguistic meaning (i.e. semantic representations, SRs) the same kind of representations that underlie non-linguistic thinking (i.e. conceptual representations, CRs)?'

To this new question, theorists divide perhaps more clearly into A and B groups. Most, but not all (e.g. Langacker 1987), of these A-theorists may simultaneously make the assumption that the 'language of thought' (i.e. CR) is universal, being part of our biological endowment. This combination of assumptions is in fact problematic, given the nature of semantic diversity, as will become clear. But first let us attend to the question whether or not there could be an identity between SR and CR, and in particular to whether or not the semantic representation of a particular linguistic string could be identical to the corresponding non-linguistic thought.

First, let us acknowledge that there might be a perhaps trivial way in which we could define a level of representation and call it CR in such a way that the elements of CR are necessarily identical to those in SR. The trivial argument would be: anything we say, we think; ergo by definition SR is identical to, or a subset of, CR. Less trivially, we may perhaps compute over SRs for the purposes of speech production, toying with alternative ways of putting things (see Levelt 1989); since these SRs might yet lack their lexical clothing, as it were, why not call them CRs? Still, that is not what is at stake. What is at stake is whether the representations we do serious thinking in (memorizing, inferring, reasoning, deciding, hypothesizing, etc.) are isomorphic with the representations that encode linguistic meanings.

Secondly, we should note that there are clearly non-linguistic thoughts that could not correspond to linguistic meanings (Keller & Keller 1996). Cases in point would be a visual image of my bedroom, appreciation of a smell, memories of the tactile properties of shark skin, and various kinds of vague, less than propositional, thoughts. Much has been written about mental imagery (q.v. Finke 1989), a little about tactile representations (Klatsky & Lederman 1993), and something about non-propositional thoughts (Sperber & Wilson 1986). It may be objected that mental imagery is not 'thinking' in the sense of reasoning, but caution is in order: consider geometry, or the literature on mental models (Johnson-Laird 1983). We surely need to concede that there is more than one 'language of thought'. Indeed, much modern work seems to underline Fodor's (1983) thesis of the 'modularity of mind', with multiple kinds of mental representation according at least to the different sensory modes of 'input', not to mention different kinds of behavioural output (speech, gesture, grosser bodily movements). Further, while Fodor imagined a single inner 'language of thought' to which these multiple specialized representations reported, and in which **real** ratiocination is conducted, more recent speculation has multiple modes of representation at the core (Levelt 1989, Jackendoff 1992, Dennett 1991).

Therefore we must transform the question once again: 'Is there a way of thinking which mirrors the way we speak, i.e. is there at least one kind of conceptual representation identical to semantic representations, such that for an arbitrary sentence the representation of the meaning just is the representation of the corresponding thought?'

Many scholars, probably those who identify with cognitive linguistics, and all the A-theorists of course, will now answer blithely yes. But for a variety of straightforward reasons, patiently pointed out by those we will call B-theorists, they must be simply wrong:

(1) Languages have both accidental lexical gaps and more systematic missing semantic fields. Tzeltal lacks a word for blue; it has a word *yax*, glossable as 'grue', i.e. green or blue: when Tzeltal speakers say in effect 'The sky is grue' they express a definite thought using the semantically general means at their disposal, in the same way that when I say *Elizabeth is my aunt* in English this does not correspond to a vague thought indifferent to whether Elizabeth is my mother's or father's sister, or the wife of my mother's or my father's brother. English just happens to lack a word for female sibling of my mother.

It is not only the 'general vocabulary' that is subject to the vagaries of uneven linguistic coding: many languages lack causal and logical connectives of the familiar sort. Guugu Yimithirr lacks a word for 'if'; you can

only express a conditional by saying in effect 'perhaps A, perhaps B', which might mean a host of other things, including 'perhaps A and B'. There is a simple moral: generally speaking, our thought is specific, but its linguistic expression is often necessarily general, non-specific, even imprecise.

These facts may seem trivial. Perhaps, in principle, any linguistic gaps can be made good by paraphrase (but see the objection in (2) below). Thus blue can be expressed as 'grue as the sky', and maternal aunt by birth as 'sister of my mother', etc. But do we think in paraphrases? Some theorists, those who subscribe to lexical decomposition, are likely to say we do. But others will point to the computational advantages of packaged concepts (further discussed below): we don't want to unpack the notion 'logarithm' every time we use it, and the general concept 'aunt' is useful in a culture where legal axioms may specify inheritance from intestate aunts without issue, regardless of the genealogical type of auntie. In addition, some notions may not be easily expressed in paraphrase at all: there are languages without real number words, others – as mentioned – without explicit logical connectives, and below we will consider the problem posed by languages that lack notions of 'left' and 'right'. Our notion 'left' as in *The ball is to the left of the tree* presupposes a triangulation between observer, ball, and tree, with coordinates based on the body-frame of the observer. There are many languages which don't use such coordinate systems, and no ready paraphrase comes to mind. A language without a conditional construction does not imply that its speakers cannot have thoughts of the kind 'if p, q'; but they cannot exactly express just such a thought in the semantic representations available to them.[1]

There are more systematic gaps in linguistic expression. As Bishop Berkeley pointed out, different kinds of representation have intrinsic limitations: a picture cannot be non-specific about metric properties and shapes, while a linguistic description can hardly avoid being so. In the same way, there are notable failures in the adequacy of linguistic expression: we are helpless when we need to describe faces, smells, contour curves, and so on. These failures of the medium are presumably related both to the existence of other 'languages of thought', other kinds of mental representation in which we think, some more analogue than digital, and to the fact that natural language vocabularies are necessarily crudely general over possible distinctions if they are to remain of learnable size.

The moral is: SR is **at most** a subset of CR, and CR is in turn just one of our representational systems or languages of thought.[2] But there are also reasons to doubt that SR is a subset of CR, as will transpire.

(2) The existence of pragmatics guarantees that we do not say what we mean (intend). If I say *Some of my books are missing*, what I probably mean or

intend is 'some but not all of my books are missing'. I don't need to spell that out given pragmatic principles governing quantifier interpretation (Horn 1972, Levinson 1983). So the semantic representation is distinct from the corresponding conceptual representation. Of course I could have said *Some but not all of my books are missing*, but I might then have implied that there was a distinct corresponding thought, namely that it was possible or even likely that all of my books might have been missing. Similarly, I didn't spell out what the thought behind the expression 'my books' was: I may have had in mind the books I own, the books I wrote, the books I had on my bookshelf (some of which belonged to the library and to friends), etc. If I had said *Some of the books I own are missing*, I might have implied that the ones borrowed from the library are not missing, and so on. Thus not only do pragmatic principles allow us to be succinct – to say less than we think – they also make it difficult to be fully explicit, for then we invariably imply more than we intended.

For this reason, Searle's (1969:19ff) **Principle of Expressability**, which holds that in principle 'whatever can be meant can be said', is not tenable for any particular case: the more explicit I try to be, the more unintended implicatures I will generate. Some theorists argue it is incorrect in principle: suppose that the lexeme *some* means **neither** 'at least some' **nor** 'some but not all', but is rather semantically general over both interpretations; then the thought must be specific where the language **cannot** be (Atlas 1989, and this volume): there is a logical form associated with the thought, but none associated with the sentence. If Guugu Yimithirr lacks a word for 'if' and makes do with a construction semantically general over 'if A, then B', 'possibly A and B', 'possibly A and not B', etc., we would be foolish to suppose that the semantic representation – as opposed to one of the possible corresponding thoughts – encodes a conditional form. On this view, semantic representations are a genus wholly distinct from conceptual representations: not only are SRs not a subset of CRs, there is no simple correspondence between them. Whereas our thought, hopefully, does indeed have some logical structure, there are no grounds for supposing that **sentences** and their corresponding SRs inherently possess 'logical forms' (Strawson 1950, Atlas 1989, Kempson 1988). Much semantic theory, the argument goes, has falsely attributed the precise articulated properties of conceptual representation to the underspecified, general nature of semantic representations.

The main contribution of recent work in linguistic pragmatics has indeed been to systematically undermine the pernicious habit of identifying the semantic content of a sentence with a determinate proposition or thought. Semantic representations must be viewed as deeply inadequate mental orphans, which serve so effectively in human communication only through

the rich interpretative principles we bring to bear on their development. Semantic theory still struggles to incorporate these insights, viewing semantic specifications as schematic (Sperber & Wilson 1986, Kempson 1988) or conceptually general (Atlas 1989) or parameterized to contexts (Barwise & Perry 1983) or at least relativized to discourse (Gazdar 1979, Kamp 1981). A-theorists, who identify CR with SR, are stubborn ostriches: they typically refuse to countenance the existence of pragmatics (Langacker 1987: 154; Jackendoff 1983: 105–106, 1992: 32) and thus are wont to attribute the determinate properties of CR to SR.

Thus SRs do not encode the corresponding thoughts, they rather – if communication is successful – invoke them. For every SR there is an indefinite range of CRs that might, in various contexts, correspond. This one-to-many mapping shows that SR and CR cannot be isomorphic. The persuasive view that SRs are schematic, incomplete, or semantically general suggests that SR is not simply a subset of CR either, but a representational medium with a different vocabulary and syntax.[3]

(3) The problem of indexicals. What is the thought corresponding to *Tomorrow I will leave here*? Could the meaning just be the thought expressed? If meanings were thoughts, then your saying *Tomorrow I will leave here* in New York on 1 September would express the same thought as my saying *Tomorrow I will leave here* in Paris on 2 August. Nor does the utterance *Tomorrow I will leave here* play the same psychic role in the individual as the corresponding thought: if days later I recollect the thought, what I think is not the meaning of *Tomorrow I will leave here*, which would have me ever itinerant. Indexical representations could not play a central role in memory; but that is just the role conceptual representations must play. As Frege (1967: 24) insisted: "If someone wants to say the same today as he expressed yesterday using the word *today*, he must replace this word with *yesterday*. Although the thought is the same its verbal expression must be different ... The mere wording ... is not the complete expression of the thought."

Indexicals open up a yawning gulf between what we say and what we think. We might hold that what we think when we say that *p* is not the semantic content nor the character or the intension, but the interpretation, some characterization of the extension, the state of affairs described by *p*. Suppose the state of affairs described by my uttering *Tomorrow I will leave here* is 'Steve will leave Duisburg on 2 August'. But suppose I mistakenly think I am in Doesburg and that tomorrow is 1 August.[4] Then what I think and what I described simply do not coincide. Cashing out indexicals in terms of determinate reference does not give us their cognitive content any more than their semantic content does. Kaplan (1977) invites us to imagine

a kidnapped heiress in the trunk of a car, saying to herself *It is quiet here now* without knowing where she is or what the time is. To capture her thought we cannot cash out the indexicals by replacing them with concrete times and places.

So what do I retrieve when I retrieve from memory the thought corresponding to *Tomorrow I will leave here*? Could it be the semantical content – now identified as conceptual representation – together with the context of the utterance? If so, the 'language of thought' contains indexicals, which must be explicated by more thoughts, which might also contain more indexicals, and so on in potential infinite regress. The difficulty is that indexicals must be banished from the language of thought, while the substituting expressions must retain their cognitive immediacy: *That tiger is after me* is not the same thought as *That tiger is after Steve Levinson* – only the former will have you hopping unless you also share the thought *I am Steve Levinson* (Perry 1993: ch. 2).

The psychological puzzles raised by indexicals run deep. They have led one psychologist, George Miller (1982: 66–67), to suppose that what is retained in memory from a demonstrative expression "must be the appearance A of some individual as determined from the perspective given by a demonstration. But this, as many cognitive psychologists would prefer, is not a linguistic or propositional object of thought. It is a (selectively incomplete) mental representation of the demonstration itself."

What is clear is that the thought corresponding to a sentence with indexicals can neither be identified with the semantic representation, the character, or the intension, nor with the interpretation or a characterization of the extension. It has to be something different if it is to play the correct role in our subjective mental life. This, I take it, is a knock-down argument against identifying SR (however interpreted) with CR.[5]

(4) Speech forces a linearization of thought and the taking of perspective. Entertain the idea of a capital T. Now describe this shape to an illiterate friend. You could say: *Draw a horizontal line. Find the mid-point, draw a vertical line of equal length downwards*. Or you could say *Draw a vertical line. Now draw a horizontal line that touches the top, with equal halves on either side of the vertical line*. There are still other ways that do not depend on notions like 'vertical': *Draw a line. Now from the end point go left 90 degrees by half the length of the first line. Go back to that point. Go right 90 degrees for the same distance*. Which system of coordinates we use and in which order we trace the shape are matters that in the spoken expression of the thought have to be determinate. But our original idea of the T might be a single gestalt, or it could embody multiple coordinate systems. This linearization of information is part of the preparation for speech production and

can be empirically investigated (Levelt 1989: ch. 4): it is particularly evident in route directions, descriptions of visual stimuli, and the like.

Closely related are forced choices in perspective. If I have a blue and a red mug in front of me, I can describe one as to the left of the other, or the other as to the right of the first. I can do both, of course, but only in one order or the other. If we describe motions, it is hard to avoid the ascription of subjective viewpoints: *the boy came in* versus *the boy went in*. It seems reasonable to suppose that we can entertain thoughts indeterminate or indifferent to perspective but find that the expression of them requires a metamorphosis of the thought. At the very least, then, speaking requires a regimentation of our thoughts, a 'thinking for speaking' as Slobin (1996) has put it, which seems to imply that the level of SR cannot be the same as the level of CR. Why would we struggle to clothe our thoughts in words if skin and cloth were the very same layer?

Notice that this is a claim apparently in the opposite direction from claims made in points (1) and (2) above, where it was pointed out that there is a tendency for natural language expressions to be semantically general where our thinking is likely to be specific. That observation however is not being undermined here: here we are noting that there are also certain systematic cases where the converse observation can be made – where a language may require semantic specificity, e.g. of a kind related to the linear nature of the medium, where our conceptual grasp of a situation may not. The two observations together constitute a powerful argument for the separation of a semantic from a conceptual level.

(5) Language-specific obligatory categories. The points made in (4) suggest that in certain systematic respects our thoughts have to be realigned for expression in a linear medium in a way general across languages. But languages also enforce obligatory grammatical and semantic distinctions that from, say, an anglocentric point of view, look alien and not the kind of thing that we imagine dominates our thoughts. For example, I can think 'You are mistaken', and speak my mind directly. A French speaker might have a corresponding thought, but she'll have to add the distinction between *tu* and *vous* if she is to voice it. I can think and say *I wish I owned that house*, but a Kwakiutl speaker will have to add the distinction 'house visible/invisible to speaker', since the determiners require such distinctions. I can say *Put the two jars on the table* without worrying about the shape of the jars, but a Tzeltal speaker has different verbs 'put on' for vessels of different shapes, and numeral classifiers that force distinctions between shape too.[6]

These facts allow two immediate conclusions. First, unless we are to maintain that e.g. a Kwakiutl speaker's non-linguistic thoughts are systematically different from our thoughts (so he can't even entertain the idea of

a house neutral with respect to visibility), we will have to hold that there is some kind of regimentation, reorganization, and embellishment of the original thought into a fully specified form that matches the grammatical and semantic structure of the language. So we have once again an argument for some kind of separation of levels. Secondly, by showing that languages differ, by virtue of their semantico-grammatical structure, in the required nature of that fully specified form, we seem to have an argument of the following form: if CR is universal, SRs various, then it can't be the case that CR = SR.

It may be objected that there is a system-use confusion here: the Kwakiutl speaker is just using a different part of the very same system as, say, the French speaker. That system is a superset of conceptual categories, a subset of which surfaces in the semantic categories of Kwakiutl. Naturally, none of our A-theorists is committed to the claim that any two individuals have the same thoughts, only that they are expressed in a single system (the universal CR = SR). In response to that objection one may note: (a) if the Kwakiutl speaker is using in both language and thinking a distinct subsystem of CR, not shared e.g. by our Frenchman, CR is now an analytical bridging construct, not a **psychological** unity; (b) if CR encompasses all the distinctions in every SR, then the things we think about as cultural and scientific discoveries were there all along in all of us (more in section 1.2); (c) if in **thinking** our Kwakiutl speaker is using the whole of CR, but in **speaking** only a systematic subpart, then this does amount to a distinction between CR and SR in the sense that psychologically there must be a separation, for the purposes of speech production, between thoughts and thoughts appropriate for expression in Kwakiutl. Introspection, and just a little evidence (see e.g. Levelt 1989, Slobin 1996), suggests that the conversion between those two levels requires some effort and acquired skill.

In summary: it is possible that SRs are distinct from CRs just in the sense that SRs are a proper subset of CRs, namely strings that include the language-specific obligatory distinctions. Then English, French, and Kwakiutl speakers can entertain the same thoughts, but they will embroider or strip those thoughts, as required by the language, when they come to speak them (this is the story in Levelt 1989). But it is also possible that as a result of those acquired habits of language production they do indeed just entertain systematically different thoughts (this will be explored below). Thirdly, as argued under points (2) and (3) above, it is possible that CRs and SRs are qualitatively different kinds of representation. The first and third possibilities would allow us to preserve the assumption of the 'psychic unity' of mankind (at the level of CR), but at the cost of making at least some kind of distinction between CRs and SRs; the second would allow us to claim an

identity between possible CRs and possible SRs, but at the cost of giving up on 'psychic unity' (further explored in 1.2 below).

(6) Languages are public representations, 'languages of thought' private ones. This point is obvious but may have deep consequences. For example, public representations have to be shared through learning; therefore there are additional constraints on their structure – e.g. vocabularies have to be of small size and thus necessarily semantically general or 'efficient' (Barwise & Perry 1983). [7] Perhaps too they have to be based, in some sense, on natural categories: blood, fool's gold, and steel girders are all ferric substances, but that might not be an easy category for a child to grasp. Given the nature of the vocal–auditory channel, and the limitations on human memory, natural languages have to be parsable in very short time spans. Given the nature of the signalling medium and the environment through which the signal passes, it will be 'noisy': hence public languages require redundancy in a way that private languages presumably do not. On the other hand, since articulation is slow and thought fast, the hearer's inference will outrun the speaker's encoding rate, so repeated concepts are best replaced with anaphoric expressions. But presumably there aren't anaphors or other such beasts in the language of thought – then we would need an inner pragmatics to unravel our thoughts, now expressed in yet another layer or 'language of thought'. And so on and so forth. A powerful but parsable and learnable syntax over a small vocabulary, whose terms are semantically general or underspecific, built on categories perspicuous to the child, with redundancy built into every message, and exploitation of anaphoric possibilities – these are special constraints on an external shared symbolic system. There is no reason to think that CRs are subject to the same constraints. And at least one reason to think that they cannot be: in that direction infinite regress looms in the shape of yet further inner languages required to make our conceptual representation as explicit as our thoughts.

There are no doubt a host of further complications for the assumption that CRs are identical to SRs.[8] But let us see where these considerations already take us:

   1. CRs cannot be identical with SRs in the sense that for every CR there is a well-formed SR that is identical to it – not all thoughts are directly encodable in all languages. That possibility would seem to be ruled out by all the arguments above.

   2. SRs cannot just be a subset of CRs, because there are essential distinctions between representations for linguistic meaning and representations for thoughts. Thoughts are the sorts of things that should be constant

when successively retrieved from memory, but SRs contain indexicals and semantically general predicates and operators which change their extensional and sometimes their intensional values or 'character' on each occasion of use.

3. SRs and CRs are not only distinct kinds of representations, distinct 'languages of thought', they are also not isomorphic. That is, the CR corresponding to the SR of a sentence is, arguably, necessarily structurally distinct. The argument for this is given in (1) and (2) above: for example, expressions in SR might be indeterminate as to logical form (cf. Atlas 1989: ch. 4; Fodor 1983: 90, 135n29). Thus we cannot get from an SR to a CR just by e.g. substituting referring expressions for indexicals. We can say that not only are SR and CR distinct, they are also non-isomorphic because there is no one-to-one mapping between relations and terms in the one 'language' and the other. Nor, if the argument from indeterminacy is correct, is SR as a language a proper subset of CR.

Thus these arguments would seem to suffice to show that CRs and SRs are, in a non-trivial way, distinct kinds of representation. This is of course contrary to assumptions in many diverse quarters, including not only Fodorians, Cognitive Linguists, and those of similar views, but also most branches of computational linguistics. Still, they are straightforward arguments and hard to refute. The arguments also happen to dispose of an extreme line of relativism, in which it is held that **we think in language**, that thinking is actually subvocal speaking, 'inner speech', or the like. Of course the arguments do not rule out the possibility of subvocal rehearsal of speech, but they do seem to establish that corresponding to those inner utterances there has to be yet another level of representation, distinct from the corresponding semantic representation, that encodes or represents the thought 'expressed'.

So why the popularity of the assumption that SR = CR? The arguments for collapsing the levels of CR and SR are essentially 'economy' arguments. Thus Jackendoff (1983: chs. 1 & 6) argues that information from various sources, including both language and vision, must come together and undergo common processes: "not to treat all these phenomena uniformly would be to miss a crucial generalisation about mental computation; hence the semantic and conceptual levels must coincide" (p.19). But this argument neglects the fact that natural language is not purely a mental object: it is a public representation system, and the constraints on shared external representation systems, as spelt out above, are distinct from whatever constraints there are on internal representation systems. Linguistic meanings are not even all in the head (see Putnam 1988: ch. 1).Thus there are reasons why we may expect the two levels to diverge.

A weaker argument for the identity of SR and CR is that, as a matter of fact, it is hard to distinguish properties of one from properties of the other,

e.g. it is hard to assign information associated with lexemes to the theoretical lexicon versus the theoretical encyclopedia, "so the burden of proof would appear to lie with anyone who would claim that linguistic semantics is autonomous" (Langacker 1987: 155). But we only have to establish one such distinction to show that SR and CR are not necessarily identical, even if they are difficult to distinguish. Indexicals, anaphoric elements, pragmatically interpreted expressions all serve to emphasize the necessity for such a distinction. Other, perhaps simpler, examples would be any case where the natives can distinguish X from Y but their lexemes lump X and Y: 'grue' or 'aunt' will do. After all, knowing the condition 'X or Y' is semantical knowledge, an artful collapse of encyclopedic entries X and Y. More telling, perhaps, is how expert knowledge can have CRs and SRs diverge. Ask a native English speaker about the meaning of jade: she'll tell you it's a semi-precious green stone; she'll know she's not able with certainty to identify it; she won't know (unless she has also read her Putnam) that mineralogically there are two distinct substances involved that happen to have the same gross physical properties -- but the jeweller will, since jadeite is worth more than nephrite. When in doubt, the final resort for CRs is always in the head, but the final resort for SRs is often (like the jeweller's) around the corner.

But granting that CRs and SRs are distinct, might we not hold that in most important respects they are likely to be similar kinds of representation? Even if they are distinct, non-isomorphic structures, perhaps they are related by a partial function, an embedding relation (Bierwisch & Schreuder 1991: 33). There are certainly plausible arguments in this direction, and they have been made ably elsewhere: (a) both kinds of representation must have the recursive generativity that allows the expression of any arbitrary fact about the world; in order to do that they must have a syntax, operating over a finite vocabulary, with rules of interpretation – in other words they must be propositional or language-like (Fodor 1983, Jackendoff 1983); (b) the speed of encoding and decoding in language from thought to language and back suggests a close similarity between CRs and SRs (Fodor 1975, Levelt 1989); given this, "we should assume that the semantic representation of a sentence is as much like the surface form as we can" (Fodor, et al. 1975: 526); (c) the learnability of languages suggests that the distance between CR and SR is small and that CR structures are suggestive of SR structures (Pinker 1989). To these I will myself add a further argument below, to the effect that SRs will induce CRs that correspond in conceptual content. Against these arguments must be placed the arguments above, especially the view that "utterance interpretation is radically underspecified by linguistic meaning" (Horn 1989: 433, discussing various views). There is a further important related issue, namely the question of 'lexical decomposition', which has implications for our next question.

1.2   *The psychic unity of the species*

Given that languages vary, do all humans essentially think alike, in the same inner conceptual representation system? The thesis of lexical decomposition is crucial here. Natural languages have single lexemes that denote complex entities or properties, for example *carburettor, damask, soul, sonata, soffit, cousin, logarithm,* and so on. If we decompose them at the level of semantic representation into constituent elementary properties, in the manner of *man* = male (x) & adult (x) & human (x), then we can capture our intuitive semantic knowledge that from the sentence *John is a man* we can infer 'John is male' by the general inference rule *(p & q)* → *q.* Alternatively, we can hold that *man* is lexically unanalysed at the level of SR, but part of knowing what *man* means is knowing the specific inference rule or 'meaning postulate' *man (x)* → *male (x).* Much ink has been spilled on this issue and there are good arguments on either side (Fodor et al. 1975; Fodor 1975; Lyons 1977; Jackendoff 1983, 1992: 48ff, etc.), but the relevance here is that one's stand on this issue combines with one's stand on the relation between SRs and CRs to yield rather radical conclusions which have received rather less attention.

It is evident (if you are in the know) that there are no words (perhaps apart from recent loans) remotely resembling in meaning *carburettor, damask, tort, sonata, pediment, licence, logarithm,* etc. in e.g. Tzeltal or Guugu Yimithirr. So if you reject lexical decomposition of such words at SR and you believe that SR is coincident with CR, then it seems that you must reject the thesis of the psychic unity of mankind and adopt radical Whorfianism. The reasoning is as follows:

1 (a) Different languages encode different concepts at a lexical level.
  (b) Different lexical concepts map directly onto distinct unitary concepts at SR.
  (c) Therefore different languages have distinct SRs (SRs with distinct vocabularies).
2 SR is identical to CR.
3 Therefore, users of different languages utilize different conceptual systems.

Now, Fodor, who subscribes to (1) and something like (2) but wishes to escape (3), is forced into a desperate further conclusion: let the level of CR include every humanly attainable concept (thus avoiding 1 (c)). Since he holds that CR, the language of thought, is an innate representation system, it will follow that, unbeknownst to the Tzeltal speakers, they knew about sonatas, logarithms, etc. all along – not in the weak Socratic sense that they might in principle have combined elementary concepts to invent the concept of the sonata, but in the sense that **that** very identical complex

whole was **already** part of their lingua mentis. Most commentators have taken this to be a reductio ad absurdam: "To have given us an innate stock of notions which includes *carburettor, bureaucrat, quantum potential*, etc., as required by Fodor's Innateness Hypothesis, evolution would have had to be able to anticipate all the contingencies of future physical and cultural environments. Obviously it didn't and couldn't, do this" (Putnam 1988: 15; see also Jackendoff 1992: 50f).

Those subscribing to lexical decomposition would seem to be in a more comfortable position (cf. Jackendoff 1992: 48ff): they can claim that the notion 'sonata' is universally attainable given an innate repository of primitive ideas, but requires a cultural and linguistic stimulus to prompt its construction from those elementary concepts. The reasoning may go as follows:

1 (a) Different languages encode different concepts at a lexical level.
  (b) Different lexical concepts map into complex formulae composed from the same set of concepts forming the universal vocabulary of SR.
  (c) Therefore different languages can have identical SRs (SRs with identical vocabularies).

2 SR is identical to CR, and CR is a universal lingua mentis.

3 Therefore, users of different languages utilize the same conceptual systems. Against this view there are a number of counter-arguments; one of the most telling is the following (Fodor et al. 1975). What is the function of complex compound concepts, if they are decomposed at the level of meaning and thought? The notion of logarithm or ellipse is useful just because once grasped it doesn't need to be repetitiously unpacked. That much we know from the psychology of 'chunking' information, or **recoding**: if I ask you to remember the letters ACUTABOVEME as I spell them to you, you may have trouble, but it's easy to remember *a cut above me* (Miller 1956, Simon 1974, Fodor 1975: 148). Because our working memories are so limited, all the computations run over the contents of that restricted storage are also affected. We seem to do not only our memorizing but also our calculating, inferencing, and thinking on the molar concepts[9] not the atomic ones from which they may be composed.[10] So whatever the universal status of the atomic concepts of CR, culture-specific molar concepts at the level of SR, and presumably CR too, probably constitute a distinct representation system in which we store our thoughts and manipulate them. This has important implications for the relation between language and thought, for languages are clearly prodigal providers of culture-specific recoded, compound concepts.[11]

All this is a thicket of complex issues, but it is probably possible to construct a synthetic view that partly bridges these different positions, preserving both the notions of the underlying 'psychic unity of mankind' and

the possibility of significant semantic and conceptual differences across cultural groups. It might go along the following lines:

1 Different languages utilize distinct lexical concepts.
2 Distinct lexical concepts are represented as distinct semantic representations on a molar level.
3 On a conceptual level, such molar concepts can be – at least in their enriched, interpreted form – broken down into atomic concepts: that is perhaps how they are learnt, but not necessarily how they are used.
4 Such atomic concepts may well be universal, from common phylogenetic inheritance or common terrestrial fate.
5 Users of different languages can therefore be said to utilize (a) **different** conceptual systems at a molar or compound level but (b) the **same** system of conceptual representation at an atomic level.
6 To the (perhaps large) extent that in thinking and speaking individuals use different concepts at the molar level, there is no 'psychic unity of mankind'; to the extent that they attain molar concepts through composition from universal atomic ones, and continue to have access to that decomposition, we may indeed talk of universal representational systems.

It is important to hold in mind the consequences of these different positions. Some of them then become quite implausible and unattractive: for example, it was certainly not Fodor's intention to advocate strong determinative Whorfianism, yet the thesis of non-compositionality leads inevitably in that direction. Similarly, the claim that CR and SR are identical, together with the facts of semantic diversity across languages, implies strong Whorfianism, unless qualified by the doctrine of lexical decomposition. The doctrine of lexical decomposition, combined with the claim that CR and SR are identical, would in its turn seem to suggest we cannot think in molar concepts, and this we know is false: indeed, it seems likely that the 'chunking' of information offered in linguistic categories is precisely the prime role of language in cognition. If we can find a way to hang on to both molar and atomic concepts in conceptual representation, then to our surprise we may find that Whorfianism, the doctrine that variant SRs map into variant CRs, and its strong anti-Whorfian contrary, the doctrine (perhaps wrong on other grounds) that SR maps directly into CR and that CR is universal, are perfectly compatible doctrines.

2   Inner and outer space: linguistic and conceptual representations of space

In the first section below I shall argue that inner, private representations cannot be totally independent from social, public ones: our innermost

.

conceptualization of experience is moulded by the language-specific concepts which at a later date we may need to express them in. This is not, in the first instance, an empirical argument, it is an argument from necessity: it is an argument from the facts of linguistic diversity to the nature of the mental codings that would be necessary to support each such distinction. But the argument begins with some surprising and little-known facts (section 2.1). The argument, developed in 2.2, is simple and direct but nevertheless leads to an unexpected, and to many no doubt unwelcome, conclusion. I expect resistance. So in the following section (2.3) I shall show that not only **must** things be just so but, as a matter of empirical fact, they **are**.

It is this last point which, one hopes, might really change the way we think about these issues. Very clever scholars have gone round and round in very different circles on the abstract question of the relation between semantics and cognition. But if we could find a way of turning the issues into a matter for empirical investigation we might get a lot further a lot faster and escape the circles. The stumbling-block has been: How do we investigate the conceptual content of the 'language of thought' independently of language? This is a psychological question, of course, but much more psychology is devoted to processing questions than to questions about the nature of the representations over which the processing is operating. Thus methods for investigating the structural properties and the conceptual content of the 'language(s) of thought' are relatively primitive. Nevertheless there is a lot of relevant work in the psychology of prelinguistic beasts – work on animals and human neonates, although the methods outlined below happen to be of independent invention.

Here then is the problem. We want to empirically investigate the relation between conceptual representations and semantic representations. We therefore need to independently investigate each system and see whether – and if so how – they are related. Let us start with the question of vocabulary. We know how to investigate the semantics of lexemes in different languages: not that there is a discovery method here, of course, but we know how to collect the extensions of a term, compare the contrastive or related terms, and thus make a guess – as the child learner must do – at the intensions. But how can we investigate whether there are corresponding non-linguistic concepts in conceptual representations? Here's a method. Invent a task that, once understood, requires no verbal input and no verbal response. But it must be so designed that the non-verbal responses will reveal the conceptual categories utilized to perform it. A simple start might be a sorting task: take a sheet of metal, a sheet of leather, and a sheet of paper, cut identical but varied shapes out of each, and ask subjects to sort the pieces any way they like – what is the salient category to them, shape or

material? Yanks sort by shape, Yucatecs by substance (Lucy 1992b). A more complex example might be: give a subject a hundred marbles and thirty sticks, and let him use the sticks to help enumerate the marbles. The correspondence he establishes between sticks and marbles may indicate an underlying basis for reckoning. Would a Mayan Indian with a linguistic numeral system to base 20 do it in the same way as you or me or a Chinese abacus-user who mentally computes complements of 5 (Stigler 1984)? In our own work we've used a simple method to make clear to the analyst the spatial coordinate system subjects are using to solve non-linguistic spatial tasks without their saying a word, as will be explained below.

If we have a method for making clear the kinds of categories that people are thinking in, not when talking but when remembering states of affairs or events, or making inferences from them, then we can treat the question 'What is the relation between linguistic and conceptual categories?' as one that might drive empirical research. That would be an advance. All we need to do is look at speakers of languages of unrelated stocks and see whether the differences in semantic categories that they use are in any way related to differences, if any, in conceptual categories employed to solve non-linguistic tasks.

Let us then turn to the spatial domain. There is a long and persuasive tradition that argues for the fundamental nature of spatial categories and systems of spatial reckoning in our mental makeup; I need only remind the reader of Kant's argument, for example, that spatial concepts are fundamental intuitions that the rest of our ideas about the material world presuppose, and so are concepts that are innate or at least inevitable and universal. So here is a domain where the weight of tradition and reason would lead us to **least** expect either variation in semantic categories across languages or variation in conceptual categories across populations; even if there were to turn out to be some linguistic differences, we might expect conceptual bedrock to remain solidly universal under our feet. It turns out that we are all in for a shock; but the function of empirical research is to deliver just those temblors.

2.1   *Different systems of spatial description*

There is a large literature on spatial distinctions in language. This gives the impression that we know exactly what these distinctions are and how they might relate to universal, presumably innate, conceptual distinctions. In brief, the literature suggests that, amongst other things, an egocentric or at least anthropocentric set of coordinates – of the sort expressed in English *left, right, front, back*, etc. – is universal in language because it is central to human spatial thinking (see Clark 1973 and Miller & Johnson-Laird 1976

for elegant, persuasive expositions of this view). But this large literature is flawed: it is not based on a serious survey of spatial description in different language families, but rather has assumed that we can leap from the semantic structure of familiar European languages to the necessary structure of mental categories.

A serious survey of spatial description in the world's language families has yet to be undertaken. But collaborative direct field investigations, undertaken by myself and colleagues in the Cognitive Anthropology Research Group of the Max Planck Institute for Psycholinguistics in about twenty languages of non-Indo-European stock, shows that linguistic description can be based on fundamentally different principles in different languages (see Max Planck Institute for Psycholinguistics annual reports for 1992 and 1993). To account for the data in our sample, we will need a typology that recognizes at least three distinct abstract types of spatial coordinate system, from which languages may select restricted subsets (Levinson 1996). In brief, we need to distinguish (a) Absolute systems of coordinates (based on something like cardinal directions), (b) Relative systems of coordinates (based on anthropocentric coordinates like our 'left', 'right', 'front', 'back'), and (c) Intrinsic systems (based on what have been called the 'inherent features' of objects, as in English *at the rear of the truck*).We leave aside here the added complications of motion descriptions and the wide range of 'topological' spatial descriptions (like English *near*) which do not employ any coordinate system at all, in the sense that they do not determine angular locations.

For present purposes I would like to contrast Absolute systems of coordinates with more familiar Relative ones. An Absolute system uses fixed bearings, like our north, south, east, and west, but they may be fixed on a completely different basis and thus have no correspondence to those directions. Now consider that there are languages where there is no other system of coordinates on the horizontal plane – for example, no Relative system of linguistic expressions like our *left* and *right, front* and *back*, which express locations relative to an egocentre or viewpoint; more rarely, there may also be little use of Intrinsic coordinates and thus no way to express notions like 'to the rear of the truck', 'by the side of the house', etc. Thus instead of saying in effect 'The boy is behind the tree', in languages like this one will have to say something like 'The boy is north of the tree'; instead of 'Take the first left, then the next right', one will have to express the directions in terms like 'Take the first turning east, then the next turning north'; instead of warning you of the snake to your left, I must talk in terms of the snake to the north, etc.

Nearly all the Australian languages of Pama-Nyungan stock are good examples of such languages (see e.g. Haviland 1992 and Levinson 1992 on

the Queensland language Guugu Yimithirr). Some elementary ruminations will show that employing such a linguistic system of description must have rather deep cognitive consequences, as we will explore briefly in the next section.

2.2   *How outer representations constrain inner representations, and linguistic concepts force non-linguistic computations*

Let us consider what special ways of thinking will be required to support such Absolute spatial descriptions. First, it is clear that speakers and their interlocutors will need to be able to locate the relevant directions in the brief time allotted in the course of conversation. In effect, they must keep a constant fix on the relevant directions, regardless of their own orientation or location, and regardless of the accessibility of external cues. Worse, they will need to dead-reckon their own current location in relation to landmarks, so they can describe e.g. route directions. Some dedicated cognitive process must constantly tick away in the background, probably combining multiple cues including inertial information and information from astronomical, meteorological, biological, and landscape features. These are the cognitive overheads of such a system of linguistic description, just as knowing constantly and unconsciously which is my left side is essential to description in terms of left and right. These predicted abilities are met in reality: speakers of such languages are not easily disoriented, and many are expert dead-reckoners (see Levinson 1992).

It is not perhaps immediately obvious that certain lexical concepts could force a specialized non-linguistic computational process. Cardinal-direction terms, if they are to be used in daily life, are a clear case. But there are many other examples. Consider using a number system to base 20, as in Mayan languages. Or consider kinship terms: even in relatively simple kin-term systems, like the Dravidian ones, which rely on an underlying distinction between 'cross' and 'parallel' kin, quite a complex algorithm is required to assign a genealogical specified relationship (e.g. X is Y's mother's brother's wife's father's son) to one of those underlying categories (Beck 1972, Levinson 1977). Or consider the application of honorifics: deciding which level to use may depend not only on assessing the kind of relationship between speaker and addressee but also between addressee and referent, and may involve also many kinds of contextual variable (see e.g. Brown & Levinson 1987). In certain cases – as with medical, mechanical, or musical terminology – the proper application of terms may require specialist knowledge and judgement (Putnam's 'division of linguistic labour'). In all such domains, utilizing the appropriate lexicon requires one to run the relevant computations: in that sense we can say that the lexical

distinctions 'force' the cognitive processes that they require. It may be that one runs the relevant computations for independent reasons, either cultural or idiosyncratic, so the sense of 'force' here is that the community-wide sharing of the lexical distinctions **guarantees** that competent speakers must make the relevant computations.

So certain lexical concepts may force certain computations – a special property of cardinal-direction terms being that they force continual background computation of orientation and location. But another important possibility is that linguistic concepts can force a difference in the way in which we experience and memorize events. Here is the argument. Some concepts are intertranslatable; for example, if I judge the table to be 2 feet high, at a later date I can easily convert my estimate to the metric system, knowing that 1 foot is near enough 30 centimetres. But some concepts or estimations, though functionally near equivalents, are not interconvertible: for example, if I judge the boy to be left of the tree, I will at a later date be unable to convert that into a cardinal direction judgement unless I have independently remembered the fixed bearings of tree and boy. There's no simple conversion algorithm, like 1 foot = 30 centimetres, relating 'left' and 'north'. The notion 'the boy is north of the tree' crucially involves ancillary information: the bearings of boy and tree. In the same way, 'the boy is left of the tree' encodes ancillary information missing from the cardinal-direction conceptualization of the scene – namely the viewpoint of the observer and his orientation with regard to boy and tree. From one coding of the scene, you cannot reconstruct the other.

This simple fact has deep consequences. A design feature of language is that we may talk – admittedly not always with any great adequacy – about anything under the sun; and a design feature of humans is that they may wish to do so. Consequently, we must memorize experiences in terms that will later make them describable. If I speak a language which exclusively utilizes cardinal-direction terms instead of terms based on egocentric coordinates, it will be no good my experiencing and recollecting events in terms of spatial concepts like 'left', 'right', 'front', and 'back': that will render me speechless – I will at least be unable to describe the spatial character of any experience. For there is no way to get from an inner coding in 'left', 'right', 'front', and 'back' terms to an outer, public linguistic coding in 'north', 'south', 'east', and 'west' terms. On the other hand, if I speak a language with 'left', 'right', 'front', and 'back' terms, it will not do to remember events solely in terms of the cardinal directions of objects and motions – I must also remember my own or some other protagonist's viewpoint on the relevant scenes. In short, coding for memory and coding for speaking cannot get too far out of line if we are to communicate about our experiences.

Cross-linguistically there are no doubt many other examples of such non-convertible, or 'untranslatable', linguistic codings of the same experience. English grammar requires us to notice if there was one stick or two – we can't say *I saw stick in the yard*; Yucatec doesn't require plural marking here, and it turns out that Yucatec speakers are less likely to notice (Lucy 1992b). From 'I saw stick in the yard' we just can't reconstruct at a later date either 'I saw a stick in the yard' or 'I saw sticks in the yard'. English names for things depend crucially on shape, Yucatec ones on the substance or stuff from which they are made. As mentioned, English speakers will classify things on the grounds of shape, while Yucatec speakers will tend to classify on the grounds of the material from which they are made (Lucy 1992b). If I classify things by substance and throw away the shape information, I won't later be able to reconstruct the shape from the substance.[12]

We can distinguish different kinds of 'Whorfian' effects – different ways in which a specific language might impinge on thinking in general – according to the temporal relationship between the relevant aspects of the thinking and the speaking.[13] There can be little doubt that at the time of speaking we think in language-specific categories; the point has been made by students of speech production (Levelt 1989) and those interested in the acquisition of languages (Slobin 1996). Slobin has dubbed this 'thinking-for-speaking' and has shown in recent collaborative work how the grammatical categories of a language then induce a certain discourse style (Berman & Slobin 1994). It is also likely that such thinking-for-speaking has consequences for recall of the scene described. In that case, thinking-for-speaking will engender changes in our memories for events. What has always been much more controversial is whether the fact that our language makes certain semantic distinctions could have anticipatory effects on the ways in which we code events at the time we experience them. The arguments here would seem to show that such Whorfian effects must indeed exist – our languages require finer-grained distinctions than we might otherwise code in memory (see also Lucy 1992a, 1992b). But readers will want more than plausibility arguments. Let us turn to some more robust, empirical evidence.

2.3   *How to think in order to speak Tzeltal*

Tzeltal is a Mayan language spoken in Chiapas, Mexico; I will report here on one dialect spoken in the highlands in Tenejapa by people I will call Tenejapans. The system of spatial description and conception has been jointly under investigation by myself and Penelope Brown from a wide range of linguistic, ethnographic, and psychological angles (see e.g. Brown 1991; Brown & Levinson 1992, 1993a; Levinson & Brown 1993; Levinson 1994). Tzeltal is a language that lacks 'left', 'right', 'front', and 'back'

notions of the kind central to English spatial descriptions; instead it utilizes Absolute coordinates, together with an elaborate and rich system of Intrinsic distinctions (Levinson 1994, Brown 1994). If an assemblage is to be described with any external frame of reference, a system of cardinal directions has to be used. It is this we shall focus on here.

2.3.1   Some distinctions in Tzeltal spatial locutions

Tzeltal cardinal directions are not directly related to celestial phenomena but are derived from characteristics of the landscape: the term that corresponds to (somewhat east of) north means literally 'down' and relates to the steep drop of Tenejapan territory from an alpine southern range to a tropical northern river valley. We will gloss the term as 'downhill'; 'uphill' therefore corresponds to a southerly direction. The orthogonal directions east and west are covered by the one term which we gloss as 'across' (although the two directions can be periphrastically distinguished, of course). This underlying three-way distinction is linguistically expressed not only in abstract nominals but in various other lexical fields, like verbs of motion: to 'ascend' is, inter alia, to head south; to 'descend', to head north; and to 'traverse', to head east or west. The dependence of the linguistic system on the overall inclination of the landscape is made patent by the fact that neighbouring peoples skew their systems according to the general slope inside their own territories. Nevertheless, these are not terms for places but cardinal-direction terms motivated by, but abstracted from, landscape features. If you take a Tenejapan outside her territory, say to a local town down on the flat lands, her 'downhill' is just where it always was, centred on a fixed bearing just east of north.

These distinctions as reflected in various lexical sets are utilized not just for the descriptions of landmarks or other large-scale features or events. They are also used for the description of small-scale locations on the horizontal. For example, to distinguish items on a table, or even to indicate which leg or ear I may have hurt, cardinal directions are likely to be used. Thus they functionally replace the 'left', 'right', 'front', and 'back' vocabulary of English, as can be demonstrated across a wide range of verbal activities and tasks (Brown & Levinson 1992).

2.3.2   Some corresponding aspects of cognition

Tenejapans are able, as one would expect, to locate their cardinal directions without difficulty. If one takes subjects into an unfamiliar concrete cell without windows and asks them to point to places to which they have been, at distances of, say, 1 to 100 miles away, they can point accurately in the correct directions.[14] As already argued, linguistic communication would fail if things were otherwise.

What difference does such a system of orientation make in how Tenejapans categorize experience and code it for memory? We have carried out a series of experiments, collaboratively designed with colleagues (see Danziger 1993), to see how Tenejapans remember spatial arrays and make inferences from them in tasks designed to probe for non-linguistic cognition. Most of these tasks are based on the following simple underlying design. Subjects see a stimulus on a table (call it table 1). They are then rotated 180° and, facing another table (call it table 2), are asked to reproduce it or recognize another token of the same stimulus type or make certain extrapolations or non-verbal inferences. For example, they might see an arrow pointing to their left, or to the east, on table 1. They are then rotated and, facing table 2, are asked to make an arrow point in the same direction. If they point the arrow to their left, then that is evidence that they coded the direction of the arrow in egocentric coordinates. If they point it to their right, i.e. again to the east, then that is prima facie evidence that they coded the direction in terms of fixed bearings. European subjects of similar age and sex (Dutch speakers in our control experiments) will point the arrow leftwards, preserving egocentric coordinates. Tenejapan subjects will generally point it eastwards, preserving the fixed bearings of the stimulus and ignoring their own rotation.

We have investigated different kinds of memory, recall, and recognition using this technique. We have also investigated non-verbal inference using the following design by Eric Pederson and Bernadette Schmitt. The subject sees two objects on table 1, say a blue cone to the left of a red cylinder. Then he is rotated and sees on table 2 the same red cylinder left of a yellow cube. Now he's rotated back to table 1 and asked to place the blue cone with respect to the yellow cube. If one interprets the arrays on table 1 and table 2 as non-verbal 'premises', one can make a simple transitive inference: blue is left of red, and red is left of yellow, so blue is left of yellow. Our European subjects make this inference. But our Tenejapan subjects tend to find a different solution, one which will preserve the directions in terms of fixed bearings. The behavioural outcomes are quite different and can easily be explained in terms of the underlying coding of spatial arrays according to different coordinate systems, egocentric (Relative) or fixed-bearing (Absolute).[15]

The details are beyond the scope of this paper (see Brown & Levinson 1993b). Suffice it to say that the results over a range of tasks are all overwhelmingly in the same direction and are persuasive evidence in favour of the thesis: if the language you speak (say, Tzeltal) primarily utilizes for spatial arrays of this sort a spatial coordinate system of an Absolute kind, then that is how you will generally remember spatial arrays and directions; while if you speak a language (like Dutch) that primarily utilizes a Relative

coordinate system, then you will generally remember spatial arrays in terms of viewpoints or egocentric coordinates. The obvious explanation for this strong, reliable correlation between verbal description and non-verbal coding for memory and inference has already been given to us by first principles: if you remember any arbitrary experience in terms of some private inner code, and that code is not intertranslatable with the public linguistic code, then you will not be able to describe your experiences. Thus the public outer code induces the coding of experience in terms rich enough to support linguistic description. In the case of Dutch and Tzeltal speakers, the outer codes are conceptually distinct in the relevant domain of spatial description; consequently, the speakers of each language will tend to code experience for non-verbal purposes in different ways, each choosing a representation congruent with the language they speak.

Might there not be alternative explanations? For example, why suppose that language categorization has the causal role? Is it not possible that for ecological or other cultural reasons Tenejapans think the way they do? Then we would have a causal effect from some other factor X that induces both linguistic and non-linguistic spatial codings of an Absolute kind. Correlations are always open to interpretations where initial causes can be pushed further back. However, in this case (a) we have a direct explanation available as just given, and (b) there is some evidence against the obvious alternatives. First, general ecological or cultural explanations are undermined by the existence of other Mayan and Mesoamerican peoples without such Absolute systems.[16] Secondly, we have experimental evidence that the kind of non-linguistic coding involved is not just a matter of remembering the fixed bearings of arrays. Rather the non-linguistic coding seems to involve **categories** that have distinct arbitrary directions – just the sort of thing that could only be shared throughout a community through the agency of a shared public language. Consider for example our notion 'north': it could have been otherwise (indeed, there are three norths often dramatically divergent: magnetic, grid, and true) – and it would not have mattered a jot, provided we all agreed on its direction. Having the notion 'north' is distinct from the ability to dead-reckon, which need make no reference to a single fixed arbitrary point or quadrant, let alone to 'north', and which is an ability that might be independently prompted by ecological or cultural conditions. So the empirical question is: Do the Tenejapans, when remembering spatial arrays, utilize a set of conceptual categories that are based on the linguistic consensus which fixes particular quadrants or cardinal directions? To explore this, we can exploit the fact that Tenejapan cardinal-directions terms have a particular property: they conflate easterly and westerly directions under the same category, conceived of as 'across' the main north–south axis. So if we get the Tenejapan subjects to perform

memory and inference tasks which require north–south distinctions they do very well; but if we take the same subjects, twist them 90°, and get them to perform the same tasks now requiring east–west distinctions, they make more errors. The conceptual coding for non-linguistic tasks appears to mirror the structure of the linguistic coding.

This raises the question: What exactly are the inner codes that seem to be involved in such performance? Could the coding for memory just be linguistic – or at least purely propositional – coding? We think not, for visual imagery and precise metric properties seem also to be involved. In some recall experiments we found that Tenejapan subjects, when shown an array of six or more complex objects in non-canonical arrangement, would still under 180° rotation preserve the fixed bearings of each object and its relations to other objects. In effect they would rebuild an arrangement as if they had viewed it from the other side, preserving non-orthogonal angles and fairly precise metric distances. This suggests that an imagistic rotation may be involved and not just a propositional encoding. A Relative linguistic description is viewpoint-based: *The cup is to the left of the jug* encodes a 'viewpoint' or egocentre – from the other side of the table, the cup must be described as to the right of the jug. In this, Relative coordinates in linguistic description are congruent with our visual images: both presuppose and preserve viewpoints. Dutch subjects asked to do the same task reproduce the array so that it looks the same from their viewpoint, even though they are rotated. But Absolute linguistic descriptions are incongruent with visual images: *The cup is north of the jug* if true remains so whichever side of the table the array is viewed from. To bring language and vision back into line, the user of Absolute coordinates must relate the two visual images by mental rotation as one 3D array under a unitary description. The Tenejapan subject is then involved in not only a different propositional coding of a spatial array but also a different processing of mental imagery, mentally rotating models of 3D arrays as he himself rotates.

There are yet further levels of inner code that may be affected by the choice of outer linguistic code. Tenejapan subjects preserve fixed bearings in their gestures. For example, when telling mythical stories that relate to real locations they point unerringly to those places at the correct compass bearings. Or when asked to describe the location of landmarks in a distant locality, they will imagine themselves at the centre of that locality and point to the bearings that the landmarks would have from that point. Pilot experiments show that at least some subjects will preserve in their gestures the fixed bearings of actions seen in a cartoon viewed on a small LCD screen-when they are rotated 180° and asked to retell the story. There is evidence from other speech communities that peoples who use Absolute linguistic coding of spatial events may systematically preserve cardinal directions in

their gestures (Haviland 1992), while those who speak languages that predominantly rely on relative descriptions will tend to preserve viewer's or protagonist's left and right in their gestures (Kita in prep.). For this to be possible, the inner code that drives gesture – presumably a specialized kinaesthetic representation system – has also been systematically affected by the properties of the external representation systems (primarily language) that individuals must utilize by virtue of their belonging to particular speech communities.

We must leave further rumination on these 'cross-modal' effects – the transfer of conceptual categories in language to many other kinds of representation system – to another essay (Levinson 1996). But I should point out that the findings are curious from the point of view of those who imagine that our mental life consists of computations over different, distinct kinds of inner representation: visual, kinaesthetic, propositional, and – if distinct – linguistic. For the normal modern presumption has been that these various representations constitute a set of innate 'languages' the same for all humans and, indeed, mostly shared by our mammalian cousins far and wide. On the other hand, these findings suggest that humans from different cultural groups might talk about, think about, visualize, and gesture about the same spatial arrays utilizing distinct concepts and distinct processes. But our arguments from first principles, as exercised above, ought to diminish the surprise – they suggest that things could not be otherwise! For information from different faculties is all grist for the linguistic mill: to talk about what we see and feel, we must be able to 'see' and 'feel' in terms that we can talk about.

3   Putting it all together

Let us itemize our interim conclusions:
1 SRs can't be identical to CRs; they can't even be homomorphic. This follows from the design properties of external languages (semantic generality and indexicality) that would make them poor internal languages (where coding for memory would be a key basic function).
2 But SRs can't be enormously remote from CRs: in particular a CR of an arbitrary event should be capable of supporting an SR describing the salient properties of that event.
3 For a CR to support an SR in this sense, the concepts available in the vocabulary of CR should be intertranslatable with the concepts available in SR. If CR has 'left' and SR only 'north', the two will not talk to one another, memories will be unretrievable or uncodable in language, and the speaker will have nothing to talk about!
4 We can empirically investigate SRs; and we can independently investigate

CRs. So the relation between SRs and CRs is a matter for empirical investigation, not for stipulation, as has been the fashion.

Where such empirical investigations will lead us we cannot now know. But preliminary findings from the cross-cultural project on spatial conception conducted by myself and colleagues suggests that most of the existing preconceptions will prove to be at best oversimplifications and most likely downright wrong.

NOTES

I have been lucky enough to have had numerous comments on drafts of this chapter from Sotaro Kita, Pim Levelt, Jan Nuyts, Eric Pederson, Jan-Peter de Ruiter, Pieter Seuren, David Wilkins, and other colleagues, together with a detailed commentary from Wolfgang Klein; whether these have been sufficient to save me from intellectual error one may doubt. Some theoretical points I now realize are less than clear, but it would take further papers to unravel them. What I hope is that there is sufficient clarity to provoke.

1 Why not? Why not just increase SR to accommodate the conceptual bulge? Because SRs are part of a public code (see point (6) below), and it takes more than my private thoughts to establish a new linguistic expression; it takes an agreement, a convention, between at least a pair of interlocutors, and more generally within an entire speech community. And as we all know, the most recalcitrant stuff in the universe is our mass of fellow beings.

2 It is important, if not easy, to systematically distinguish the following different kinds of way in which representation systems might be identical or distinct: (a) a token string in one system might be an identical string of symbols to a token in another system (here too we should worry about whether 'string' implies labelled bracketing or other indications of disambiguated parsing); (b) all the strings generated by one system might be strings in the other system, but all the strings in the latter might or might not be strings in the former – i.e. one system might be a subset, proper or otherwise, of the other; (c) all the strings in one system might be distinct from all the strings in the other, but for each string in one there might be a string in the other with identical interpretation, or not, as the case may be; (d) etc. It takes considerable care to keep these all apart, and whether it is worthwhile or not will depend on theoretical motivations. Since there are developed techniques for talking about all this in mathematical linguistics (see e.g. Partee, ter Meulen & Wall 1990), our reluctance to employ them is perhaps indicative of an underlying embarrassment associated with treating the 'language of thought' hypothesis as something more serious than metaphor.

3 And syntax? Yes, if the arguments of Atlas (1989) and Kempson (1988) are right: then the logical character of our thought, such as it is, might be captured in some canonical logical form, absent from semantic representations, "for ordinary language has no exact logic" (Strawson 1950: 344).

4 Duisburg and Doesburg are towns on either side of the Dutch–German frontier.

5 At least it establishes that the vocabularies of SR and CR are distinct; the pragmatic arguments mentioned in (2) above seek to establish that the syntax is also distinct, by showing that logical form is not a property of SR but – it may be

supposed – of CR. Wolfgang Klein points out to me that these two kinds of arguments are more parallel than the presentation suggests. One might for example put it as follows: just as CR must lack indexicals while universally all SRs have them, so CR must have logical variables while languages and their SRs universally lack a direct encoding of them (some current views notwithstanding – here see Kempson 1988).

6 For the full Tzeltal facts see Berlin 1968 (on numeral classifiers), Brown 1994 (on locative predicates), Levinson 1994 (on body parts as locative specifications).

7 This itself is a telling oversimplification: a concept in a public representation doesn't need to be fully shared – for example, I can use the word *manslaughter* even though I am less than sure of the legal definition, because we know that the lawyer round the corner could if necessary set us straight. From this point – that meaning can be distributed in a social network – much follows about why meaning can't be purely in the head (see Putnam 1988).

8 One such difficulty is that there are a range of neurological deficits that seem to show selective impairment of just one of these representations. For an interesting case, see the patient described in Hart & Gordon 1992, who had "cerebral damage whose pattern of acquired deficits offers direct evidence for a major division between visually based and language-based higher-level representations" (p. 60). The selective deficit was in language concerning the physical properties of animals, which were readily apprehended visually. The issues are however not entirely straightforward: see Caramazza, Hillis, Rapp & Romani 1990. (I am grateful to Pim Levelt for pointing this literature out to me.)

9 Incidentally, the traditional opposition is between 'molar' and 'molecular' levels, but partly to avoid alliterative confusion, and partly to employ the more familiar phrase 'atomic concept', I'll oppose 'molar' and 'atomic'.

10 Miller's (1956) original insight was that the unit for working memory is such a 'chunk' of familiar material, and that working memory – over which our computations operate – is restricted to about six chunks (and not, as presupposed by the information theory of the time, to so many BITS of information). It is now known that there are additional absolute temporal constraints on working memories of specific kinds, and that the two kinds of constraints together seem to predict the patterns of human performance (Baddeley 1990: 81–85). There is also a working hypothesis that long-term memory is constrained by chunking, so that each chunk – however composed internally – takes, say, 8 sec to store (Simon 1986: 303).

11 Fodor (1992: 389), for one, seems willing to buy these implications:

So I am not committed to asserting that an articulate organism has no cognitive advantage over an inarticulate one. Nor, for that matter, is there any need to deny the Whorfian point that the kinds of concepts one has may be profoundly determined by the character of the natural language one speaks. Just as it is necessary to distinguish the concepts that can be expressed in the internal code from the concepts that can be entertained by a memory-restricted system that computes with the code, so, too, it is necessary to distinguish the concepts that can be entertained (*salve* the memory) from the ones that actually get employed. This latter class is obviously sensitive to the particular experiences of the code user, and there is no principled reason why the experiences involved in learning a natural language should not have a specially deep effect in determining how the resources of the inner language are exploited.

12 I would like to be able to be more precise here about such 'untranslatable' codings of the same state of affairs. Perhaps something along the following lines might do.

   We could invent a term for concepts – let us call them **intransformables** – which are closely related extensionally but intensionally distinct. Informally, an intransformable pair of expressions are two expressions which are 'about the same domain' but which contain ancillary information, so that one cannot get from the one expression to the other. A bit more precisely:

   Let us say that a and b are intransformables if a is part of a system of concepts A and b part of a system of concepts B, such that A and B cover the same extensional domain, but (for arbitrary sentence frame S) Sa and Sb are not analytically related.

   We might want to distinguish **unilateral intransformables** where, in a suitable sentence frame S, Sa entails Sb, but Sb does not entail Sa, from **bilateral intransformables** where no entailment relations hold between Sa and Sb. Then *grue* is a unilateral intransformable with respect to *green*, because from *This chair is grue* I can't infer 'this chair is green', even though from *This chair is green* I may infer 'this chair is grue'. On the other hand, *left* is a bilateral intransformable with respect to *north*: *The boy is north of the tree* and *The boy is left of the tree* can both be true of the same state of affairs, yet we cannot infer one description from the other. Moreover, despite the conceptual difference, the two descriptions could play the same functional role: if we hold the viewpoint constant, and ensure that north is aligned with left, then the set of possible worlds in which one description is true is identical to the set in which the other is true, and similarly for *right* and *south*, etc. It is this systematic functional and extensional near-equivalence together with intensional distinctiveness that we are trying to capture.

13 This paragraph records thoughts emerging from some intense seminar discussions at the Max Planck Institute for Psycholinguistics in 1993 with Pim Levelt, Dan Slobin, John Lucy, and others.

14 There was extremely high consensus, or small angular deviation, across subjects; actual accuracy was clearly affected by the orientation of the room, which was on the diagonal of the 'uphill'/'downhill' axis. This itself is interesting, because it shows – a point made below – that categorical information is involved in their thinking about relative locations in a non-linguistic task.

15 The terms 'egocentric' and 'Relative' are not the same concepts, although here they happen to equate; see Levinson (1996) for exposition, which would take us far afield here.

16 Colleagues have investigated some of these other cultures. For example, Eve Danziger reports that Mopan speakers in Belize, who speak a Mayan language and share the basic Mesoamerican maize-culture complex, utilize no Absolute system in the same kinds of tasks.

REFERENCES

Atlas, Jay. 1989. *Philosophy without ambiguity*. Oxford: Clarendon Press.
Baddeley, Alan. 1990. *Human memory*. Hillsdale, N.J.: Erlbaum.
Barwise, Jon, & John Perry. 1983. *Situations and attitudes*. Cambridge: MIT Press.

Beck, Brenda. 1972. *Peasant society in Konku: a study of right and left subcastes in South India.* Vancouver: University of British Columbia Press.

Berlin, Brent. 1968. *Tzeltal numeral classifiers: a study in ethnographic semantics.* The Hague: Mouton.

Berman, Ruth, & Dan Slobin. 1994. *Relating events in narrative: a crosslinguistic developmental study.* Hillsdale, N.J.: Erlbaum.

Bierwisch, Manfred, & Rob Schreuder. 1991. From concepts to lexical items. In W. Levelt (ed.), *Lexical access in speech production.* Oxford: Blackwell, 23–60.

Brown, Penelope. 1991. Spatial conceptualization in Tzeltal. Working Paper No. 6. Cognitive Anthropology Research Group, Nijmegen.
  1994. The INs and ONs of Tzeltal locative expressions: the semantics of static descriptions of location. In J. B. Haviland & S. Levinson (eds.), special issue of *Linguistics* 32: 743–790.

Brown, Penelope, & Stephen Levinson. 1987. *Politeness.* Cambridge: Cambridge University Press.
  1992. Left and right in Tenejapa: investigating a linguistic and conceptual gap. *Zeitschrift für Phonetik, Sprachwissenschaft und Kommunikationsforschung* 45(6): 590–611.
  1993a. 'Uphill' and 'downhill' in Tzeltal. *Journal of Linguistic Anthropology* 3(1): 46–74.
  1993b. Linguistic and nonlinguistic coding of spatial arrays: explorations in Mayan cognition. Working Paper No. 24. Cognitive Anthropology Research Group, Nijmegen.

Caramazza, A., A. E. Hillis, B. C. Rapp & C. Romani. 1990. The multiple semantics hypothesis: multiple confusions? *Cognitive Neuropsychology* 7(3): 161–189.

Clark, Herbert. 1973. Space, time, semantics and the child. In T. E. Moore (ed.), *Cognitive development and the acquisition of language.* New York: Academic Press, 28–64.

Danziger, Eve (ed.). 1993. Manual for the field administration of cognitive spatial tasks. Cognitive Anthropology Research Group, Nijmegen.

Dennett, Daniel. 1991. *Consciousness explained.* Boston: Little, Brown.

Finke, Ronald A. 1989. *Principles of mental imagery.* Cambridge: MIT Press.

Fodor, Janet, Jerry Fodor & Merrill Garrett. 1975. The unreality of semantic representations. *Linguistic Inquiry* 4: 515–531.

Fodor, Jerry. 1975. *The language of thought.* Cambridge: Harvard University Press.
  1983. *The modularity of mind.* Cambridge: Harvard University Press.
  1992. How there could be a private language. Extract from Fodor 1975, reprinted in B. Beakley & P. Ludlow (eds.), *The philosophy of mind.* Cambridge: MIT Press, 385–391.

Frege, Gottlob. 1967. The thought: a logical inquiry. In P. Strawson (ed.), *Philosophical logic.* Oxford: Oxford University Press, 17–38. (Originally published 1918)

Gazdar, Gerald. 1979. *Pragmatics.* New York: Academic Press.

Hart, J., & B. Gordon. 1992. Neural subsystems for object knowledge. *Nature* 359: 60–64.

Haviland, John. 1992. Anchoring, iconicity, and orientation in Guugu Yimithirr pointing gestures. *Journal of Linguistic Anthropology* 3(1): 3–45.

Horn, Laurence. 1972. *On the semantic properties of logical operators in English.*

Indiana University Linguistics Club, Bloomington. Mimeo.

1989. *A natural history of negation*. Chicago: University of Chicago Press.

Jackendoff, Ray. 1983. *Semantics and cognition*. Cambridge: MIT Press.

1992. *Languages of the mind*. Cambridge: MIT Press.

Johnson-Laird, Philip. 1983. *Mental models: towards a cognitive science of language, inference, and consciousness*. Cambridge: Cambridge University Press.

Kamp, Hans. 1981. A theory of truth and semantic representation. In J. Groenendijk, T. Janssen & M. Stokhof (eds.), *Formal methods in the study of language*, vol. 1. Amsterdam: Mathematisch Centrum, 277–320.

Kaplan, David. 1977. Demonstratives. Paper presented at the Symposium on Demonstratives, American Philosophical Association. Ms., University of California at Los Angeles.

Keller, Janet, & Charles Keller. 1996. Imaging in iron or thought is not inner speech. In J. J. Gumperz & S. Levinson (eds.), *Rethinking linguistic relativity*. Cambridge University Press, 115–129.

Kempson, Ruth. 1988. Grammar and conversational principles. In F. Newmeyer (ed.), *Linguistics: the Cambridge survey*, vol 2. Cambridge: Cambridge University Press, 139–163.

Kita, Sotaro. In preparation. Experiments on coding of visual stimuli in gesture. Max Planck Institute for Psycholinguistics, Nijmegen.

Klatzky, Roberta, & Susan J. Lederman. 1993. Spatial and non-spatial avenues to object recognition by the human haptic system. In N. Eilan, R. McCarthy & B. Brewer (eds.), *Spatial representation*. Oxford: Blackwell, 191–205.

Langacker, Ronald. 1987. *Foundations of cognitive grammar*, vol. 1: *Theoretical perspectives*. Stanford: Stanford University Press.

Levelt, Willem. 1989. *Speaking: from intention to articulation*. Cambridge: MIT Press.

Levinson, Stephen. 1977. *Social deixis in a Tamil village*. Ph.D. thesis. Ann Arbor Microfilms.

1983. *Pragmatics*. Cambridge: Cambridge University Press.

1992. Language and cognition: the cognitive consequences of spatial description in Guugu Yimithirr. Working Paper No. 13. Cognitive Anthropology Research Group, Nijmegen.

1994. Vision, shape and linguistic description: Tzeltal body-part terminology and object description. In J. B. Haviland & S. Levinson (eds.), special issue of *Linguistics* 32: 791–855.

1996. Frames of reference and Molyneux's question: cross-linguistic evidence. In P. Bloom, M. Peterson, L. Nadel & M. Garrett (eds.), *Language and space*. Cambridge: MIT Press, 109–169.

Levinson, Stephen, & Penelope Brown. 1993. Immanuel Kant among the Tenejapans: anthropology as empirical philosophy. *Ethos* 22(1): 3–41.

Lucy, John. 1992a. *Language diversity and thought*. Cambridge: Cambridge University Press.

1992b. *Grammatical categories and thought*. Cambridge: Cambridge University Press.

Lyons, John. 1977. *Semantics*. 2 vols. Cambridge: Cambridge University Press.

Miller, George. 1956. The magical number seven, plus or minus two. *Psychological Review* 63(2): 81–97.

1982. Some problems in the theory of demonstrative reference. In R. Jarvella & W. Klein (eds.), *Speech, place and action*. New York: Wiley, 61–72.

Miller, George, & Philip Johnson-Laird. 1976. *Language and perception*. Cambridge: Cambridge University Press.

Partee, Barbara, Alice ter Meulen & Robert E. Wall. 1990. *Mathematical methods in linguistics*. Dordrecht: Kluwer.

Perry, John. 1993. *The problem of the essential indexical*. Oxford: Oxford University Press.

Pinker, Steven. 1989. *Learnability and cognition*. Cambridge: MIT Press.
1994. *The language instinct*. New York: Morrow.

Putnam, Hillary. 1988. *Representation and reality*. Cambridge: MIT Press.

Searle, John. 1969. *Speech acts*. Cambridge: Cambridge University Press.

Simon, Herbert A. 1974. How big is a chunk? *Science* 183: 482–488.
1986. The parameters of human memory. In F. Klix & H. Hagendorf (eds.), *Human memory and cognitive capabilities*. Amsterdam: North Holland, 299–310.

Slobin, Dan. 1996. From 'thought and language' to 'thinking for speaking'. In J. J. Gumperz & S. C. Levinson (eds.), *Rethinking linguistic relativity*. Cambridge University Press, 70–96.

Sperber, Dan, & Deirdre Wilson. 1986. *Relevance*. Oxford: Blackwell.

Sterelny, Kim. 1990. *The representational theory of mind*. Oxford: Blackwell.

Stigler, James. 1984. Mental abacus: the effect of abacus training on Chinese children's mental calculation. *Cognitive Psychology* 16: 145–176.

Strawson, Peter F. 1950. On referring. *Mind* 59: 320–344.