

KENT JOHNSON

ON THE NATURE OF REVERSE COMPOSITIONALITY

ABSTRACT. Reverse Compositionality (RC) is the thesis that one understands a complex expression only if one understands its parts. I argue that this thesis is false for natural languages. I then argue that the phenomenon that motivates the thesis is more likely to be a fact about human sentence-processing than linguistic understanding per se. Finally, I argue that RC is not useful in the debates about prototype-style theories of concepts in which it figures heavily.

1. INTRODUCTION: A FIRST PASS AT REVERSE COMPOSITIONALITY AND ITS INTEREST

Over the years philosophers have uncovered a number of principles regarding how language is structured. First, perhaps, was the principle of Compositionality, which says that the meaning of a complex expression is a function of the meanings of its parts and their mode of combination (e.g. Frege, 1892). Roughly speaking, if you understand the meaning of ‘dogs’ and the meaning of ‘barks’, and you know what is expressed by a subject-predicate sentence of the relevant form, then you have enough information to figure out the meaning of the sentence ‘Dogs bark’. We also have the principle of Systematicity, which says roughly that expressions of the same linguistic type can be substituted for each other *salve grammaticalitate* (e.g., Fodor and Pylyshyn, 1988). Thus, if you understand the sentences ‘Dogs bark’ and ‘Cats meow’, then you also understand the sentences ‘Dogs meow’ and ‘Cats bark’. Recently, a third principle of language has entered the scene: Reverse Compositionality (RC) (e.g., Fodor, 1998a, b; Fodor and Lepore, 2001; Dietrich and Markman, 2003, pp. 109–110; Robbins forthcoming). RC is similar to Compositionality in that it concerns the relationship between the meaning of a complex expression and the linguistically primitive expressions and syntactic structures it is composed of. But instead of saying that you can determine the meaning of the complex by determining the meanings of its parts, RC says that you can determine the meaning of

the parts by determining the meaning of the complex. For instance, if you understand the meaning of a complex expression like ‘Dogs bark’, then you understand the meaning of simple expressions like ‘dogs’ and ‘bark’. Moreover, this relation between understanding ‘dogs bark’ and understanding ‘dogs’ and ‘bark’ isn’t an accident. As Fodor writes, ‘it is something like true by definition that mastering a complex expression requires mastering its constituents, since, after all, constituents are by definition *parts* of their hosts’ (Fodor 1998b, p. 52). Given this characterization, we might, as a first pass, characterize RC as follows:

- (1) A person understands the meaning of a (non-idiomatic)¹ complex expression of her language only if she understands the meanings of the primitive expressions and syntactic configurations that the former is built out of.

The characterization of RC given in (1) will loom large in what is to follow, and I’ll have more to say about it later. At this point, though, it will be useful to see what work RC is supposed to do.

The primary use for RC is to defend a certain view about the nature of thought. This is done by assuming that if RC holds for language, then it should hold for thought, too (cf. e.g. Fodor 1998b, p. 49, n. 1) and Fodor and Lepore (2001, p. 351, n. 1). (I’ll have more to say later about Fodor’s tendency to “move back and forth between talking about concepts and talking about words as ease of exposition suggests”). Fodor maintains that RC shows that much current research in cognitive science about concepts is wrong-headed. RC, he argues, supplies the crucial premise for the claim that concepts cannot be prototypes, stereotypes, exemplars, ‘recognitionally’ based, etc. (Fodor, 1998a, b). In fact, Fodor argues that RC shows that a concept cannot express more than what the concept contributes to the complex concepts in which it occurs. The argument for this last claim goes as follows. Suppose that a concept C did express some ‘extra’ feature, which is not contributed to the complex concepts in which C occurs. Then it would be possible for a person to possess the complex concept without possessing C. But this outcome is tantamount to saying that one could understand the expression ‘pet fish’ without understanding ‘pet’, or ‘fish’. But by RC, Fodor argues, this is impossible:

it is something like true by definition that mastering a complex expression requires mastering its constituents, since, after all, constituents are by definition parts of their hosts. So, to say that you could master ‘pet fish’ without mastering ‘pet’ (or, *mutatis*

mutandis, 'red triangle' without mastering 'red') is tantamount to saying that 'pet' isn't *really* a constituent of 'pet fish' after all; which is, in turn, tantamount to saying that 'pet fish' is an idiom. Which, however, 'pet fish' patently is not. So I win. (Fodor, 1998b, p. 52)

The argument just given is meant to be highly destructive, since all the types of concepts listed above treat concepts as having features which may not be expressed by complex concepts that contain the concepts. For instance, a prototype theory might include in the concept PET the properties of being furry and cuddly, whereas neither of these properties are part of the prototype of the complex concept PET FISH. Thus, by the argument just given, prototype theories are false.

Clearly, RC plays a central role in this argument of Fodor's, and I'll return to it later. However, my primary interest in this paper is to explore RC in its own right. Although RC is sometimes taken as a claim about thought, not language, I have characterized it as a claim about language because, scientifically speaking, there is very little of relevance that can be uncontroversially maintained about thought, whereas language is a bit easier to look at directly. (I will discuss RC for thought later, though.) RC engenders some strong but conflicting intuitions among philosophers. In discussing the issue, I have encountered numerous philosophers who have asserted the (trivial) falsity of claims like (1) with as much vigor and enthusiasm as Fodor asserts its (trivial) truth. But both these camps are wrong. On its most natural and interesting interpretation, (1) is neither obviously true nor obviously false. Indeed, Fodor himself seems to have noticed that RC is not obviously true: as Robbins (forthcoming) discusses, Fodor actually offers a number of arguments in support of the truth of RC. Rather, (1) is a significant empirical claim about human cognitive and linguistic abilities. Later, I'll argue that (1) is not always true. Let me now say why it's not obviously false.

(1) might seem questionable if we interpret it as some sort of conceptual or metaphysical truth. (Indeed, Fodor's comments about RC being true 'by definition' suggest as much.) But then (1) is rather uninteresting. Of course, it's *possible* that (1) would fail to hold: it's all too easy to describe possible worlds in which agents understand complex expressions without understanding their constituents. Claims like (1) are much more interesting if they are interpreted with a weaker modality, as reflecting only a general psychological law. In that case, it will not be sufficient to simply show that (1) could – in the barest sense – admit of counterexamples. Rather, it's interesting to see whether (1) holds in utterly *normal* circumstances. In such a case,

the only relevant counterexamples to (1) will be those that involve *plausible* circumstances. This limitation on the range of acceptable counterexamples makes the falsification of (1) considerably more difficult. After all, it's just not true that normal people in normal circumstances understand 'pet fish' without understanding 'pet' or 'fish' !

A word or two about the notion of understanding is also in order. One might try to make (1) a conceptual truth by redefining 'understanding' so that understanding a complex entails understanding its constituents. But there are problems with such a maneuver. In particular, it appears to be a genuine redefinition of 'understanding'. It's easy to imagine that Shelia has a complete and total (apparent) understanding of some complex expression ABC, even though she doesn't understand some of its constituents. Shelia can use ABC in conversations, and associates it with the appropriate concept(s) whenever she hears others use it. She's even aware that ABC is composed of constituents, but simply doesn't understand at least one of them. For a simple image of how this could happen, suppose that the meaning of a propositional logic sentence is simply its truth-value. It would then be easy to understand ' $(P \rightarrow Q)$ ' without understanding 'P'. (Better yet, one could know that $\$(\#P @ \$Q)$ expresses the material conditional without knowing whether @ expresses conjunction or disjunction (or whether \$ or # express negation or the constant function).) This sort of phenomenon also occurs in more empirically plausible situations using actual natural language (cf. see Section 2 for detailed discussion of some examples). Moreover, protecting (1) from empirical scrutiny by such a redefinition of the term doesn't resolve any issues, but merely relabels them. If Shelia's complete and total competence with ABC (albeit not with all of its constituents) doesn't count as understanding, then it's clear that there's an important cognitive relation she bears to ABC which is crucial to a theory of human linguistic abilities. Thus, throughout this paper, I'll use 'understanding' in a fairly intuitive sense. Roughly speaking, I'll say that a person understands an expression if she possesses and assigns the right concept to it.

My goal in this paper is to sort out what is correct and what is incorrect about RC as it applies to language. I will argue that there is a fundamental flaw in (1). Thus, like Compositionality and Systematicity, RC stands in need of clarification.² In Section 2, I show that (1) is sometimes false in ordinary circumstances. Moreover, the general structure of the counterexample appears in many places. In Section 3, I turn to the question of what is right about

RC. I'll address this question by replacing (1) with a different theory of RC. For the bulk of this paper, I will treat RC as an epistemic thesis. However, RC is sometimes expressed as a metaphysical thesis, stating roughly that the meanings of a complex expression's parts is *determined* by the meanings of the complex. In Section 4, I examine this metaphysical thesis, and argue that there is little to be said for it. In Section 5, I will revisit Fodor's argument against prototype (and exemplar and recognitional, etc.) theories of concepts. We'll see that, contrary to Fodor's assertions, the argument does not wreak havoc on the theories of concepts in its scope. I conclude in Section 6.

2. THE ANATOMY OF A COUNTEREXAMPLE TO REVERSE COMPOSITIONALITY

(1) is the most natural characterization of RC. Unfortunately, (1) is false. To see why this is so, I'll use some ideas that Zeno Vendler made famous in philosophy and linguistics (Vendler, 1967, ch. 4). To begin, consider a verb like 'build', as in (2)

(2) Mary built the house.

Semantically speaking, 'build' in (2) picks out a completed event of building, where this event has a temporal ending at the point at which the house is built.³ Following Vendler, I will call verbs that specify a temporal endpoint *telic* verbs. Thus, on a natural interpretation of 'Mary built the house in two weeks', the prepositional phrase 'in two weeks' picks out the duration of the building event, which ended when the house became built. But now consider a sentence like (3)

(3) Mary was building the house.

Unlike (2), (3) does not pick out a completed event of building the house. For instance, (3) is a natural part of the sentence 'Mary was building the house when she died, and so the house was never built'. Semantically speaking, the progressive morpheme '-ing' removes the specification of an endpoint in the meaning of 'build', leaving (3) with only the meaning that Mary was engaged in a house-building activity (which may or may not have reached completion). We can see this when we note the oddity of the sentence *'Mary was building the house in two weeks'. Although this sentence has an irrelevant meaning according to which Mary started building the house within two weeks of some contextually specified time, the phrase 'in two

weeks' can no longer pick out how long the building went on before the house was completed, as it did in (2).

The contrast between (2) and (3), I claim, provides a straightforward counterexample to RC. For a person could understand the meaning of (3), and yet fail to understand the meaning of 'build'. For instance, he might incorrectly interpret 'build' as only expressing that some kind of building activity is occurring, without any specification of an endpoint to this activity. Such a person would interpret 'Mary built the house' as true in a situation where Mary laid the foundations and put up only one of the walls of a house. Thus, he might interpret 'build' as an *atelic* verb like 'watch', which does not specify an endpoint of the activity. Just as the progressive morpheme can apply to 'watch' – e.g. 'Mary watched the kitten', 'Mary was watching the kitten' – without removing any specification of an endpoint, it could also apply to the present misinterpretation of 'build'.⁴ Errors involving the confusion of telic with atelic expressions occasionally happen when people are learning the expression. Even worse, a person could understand (3) without having any idea whether 'build' was telic or atelic. In this case, he would not have an incorrect view about the telicity (and hence the meaning) of 'build', but would have no view at all. Such a person could have a very clear understanding of 'Mary was building the house', even though he wondered whether 'build' is telic or atelic. (In fact, he might even mistakenly think that there are two words 'build', one telic and the other atelic.) In short, contrary to what (1) predicts, a person could understand the meaning of a complex verb phrase like 'building the house' without understanding one of its constituents, namely 'build'.⁵ So (1) is false.

The above counterexample certainly meets our requirement of empirical plausibility. In fact, I have experienced something like it firsthand. I worked in a factory one summer, and when I was being trained, I was briefly shown a process whereby the parts that we manufactured were put into a machine which coarsened them. I was told that this was where you 'flaked' the parts, and that for a big order, you could be 'flaking' parts all night. I didn't know whether the verb 'flake' was telic or atelic, although I fully understood such claims as 'Kevin is flaking the parts tonight'. (According to the machine's operating instructions, 'flake' is atelic: it involves bombarding an object with small metal flakes; the fact that some but not all of the parts look flaky afterwards has nothing to do with the verb's meaning.)

Before moving on, let me address five questions about the above type of counterexample. First, could it be that the difference

between telic and atelic verbs is ‘too small’ for the above argument to constitute a genuine counterexample to (1)? It’s hard to see how this could be so, since a verb’s telicity is widely regarded as one of its linguistically most important semantic features (e.g., Vendler, 1967; Verkuyl, 1989; Parsons, 1990; Tenny, 1994; Higginbotham, 2000; and many others). Additionally, the telic/atelic distinction has played a substantial role in philosophy since ancient times. Aristotle was keenly aware of it. In fact, this distinction is at the heart of his famous distinction between movement (*kinesis*) and actuality (*energeia*). He writes:

Every movement is incomplete – making thin, learning, walking, building; these are movements, and incomplete movements. For it is not true that at the same time we are walking and have walked, or are building and have built, or are coming to be and have come to be . . . but it is the same thing that at the same time has seen and is seeing, or is thinking and has thought (*Metaphysics* 1048b 18–34; cf. 1049a1–4, 193a31–b6)

Besides being widely recognized, in formal semantics the telic/atelic distinction is explicitly marked in the meanings of sentences. Parsons (1990), along with many others, treats this distinction as a difference in the theoretically primitive conjuncts contributed to the meaning of a sentence. In Parsons’ theory, atelic verbs express that the action of the verb ‘holds’ or continues on in time, whereas telic verbs express that the action ‘culminates’. (The underlying logic behind these notions is built into every current theory of tense and aspect; I use Parsons’ theory as an example only because of its relative simplicity and perspicuity.) Thus, the meaning of ‘Mary watched the kitten’ and ‘Mary was watching the kitten’ is given by (4a), whereas the meaning of ‘Mary built the house’ is given by (4b), and the meaning of ‘Mary is building the house’ is given by (4c) (I have ignored the tenses to reduce clutter.):

- (4) a. $\exists e(\text{Watch}(\text{Mary}, \text{the kitten}, e) \wedge \text{Hold}(e))$
 b. $\exists e(\text{Build}(\text{Mary}, \text{the house}, e) \wedge \text{Culminate}(e))$
 c. $\exists e(\text{Build}(\text{Mary}, \text{the house}, e) \wedge \text{Hold}(e))$

Thus, Parsons explicitly represents telicity and atelicity by the predicates ‘Hold’ and ‘Culminate’, and the effect of the progressive is to replace any occurrences of ‘Culminate’ with occurrences of ‘Hold’. The representation of telicity and atelicity is a standard component of a great many other semantic theories, as well (e.g., Verkuyl, 1989; Landman, 1992; Tenny, 1994; Bonomi, 1997; Higginbotham, 2000). Thus in terms of both its notoriety and its

impact on formal semantics, the telic/atelic distinction is clearly substantial.⁶

As an aside, one might worry that appealing to notions such as telicity is unfair to Fodor and Lepore, who elsewhere deny the existence of semantic structure in typical words (e.g. Fodor and Lepore, 1998). But the explanations I have offered are almost universally accepted in the linguistics community, and they accommodate a broad range of phenomena across hundreds of typologically unrelated languages. Fodor and Lepore suggest that an expression's telicity is not a semantic feature, but is only a pragmatic phenomenon, explained by our worldly knowledge. Thus, the fact that 'Mary built the house' implies a certain end-point to the activity is due simply to our knowing about what it is to have built something. But our pragmatic abilities don't explain why most speakers find the sentences in (**) odd, since they both have perfectly clear meanings

- (**)
- a. *John ran a mile for ten minutes [Meaning: There was a period that lasted ten minutes, and in that period, John ran a mile.]
 - b. *Amy wrote books in ten years [Meaning: there was a period in Amy's life during which she wrote books, and this period ended after exactly ten years.]

The judgments in (**) are typically explained by reference to linguistic properties such as telicity. More generally, there is a massive amount of evidence in the linguistics literature that properties like telicity play a major linguistic role in the organization of a clause. (E.g., in many languages, e.g., Greek and Bulgarian, the telicity of a verb is morphologically marked, much as the progressive '-ing' can affix some verbs in English. Such markings are hard to deny, regardless of where and how in a linguistic theory one chooses to implement verbal aspect.) The systematic patterns in this evidence – often appearing overtly in the sentence structure – is what makes linguists tend to think that telicity is a genuinely linguistic phenomenon. Until Fodor and Lepore can supply some inkling of how to account for these many phenomena in a theory that is at least as good as our current ones, it's hard to see how their denial can have much force. In any case, we will see that the example offered above can be reproduced with many other constructions.

Second, could there be something problematic about progressive constructions, so that the counterexample doesn't really work? Maybe the counterexample is somehow suspect because '-ing' is just

a lowly morpheme that attaches to a word. Perhaps ‘build’ and ‘building’ are altogether different words, so the fact that a person could understand one but not the other is no threat to (1). Thus, maybe understanding (3) but not (2) is no different than understanding ‘Mary painted the house’ but not ‘Mary cleaned the house’. The problem with this move is that the class of progressive verb phrases is productive in two distinct ways. In the first place, it is universally accepted in theoretical linguistics that the syntactic structure of a progressive verb phrase that is relevant to interpretation has the form: Tense[Progressive[VP]] (e.g., Parsons, 1990; Drijkoningen, 1992; Landman, 1992). So although the progressive morpheme ‘-ing’ attaches to an individual verb, it actually functions semantically as an individual word, taking scope over the entire remaining verb phrase. (In this sense, the progressive is similar to the tense inflection on a verb.) So in terms of interpretation, the verb phrase in (3) is revealed by the structure: Past[Progressive[‘build a house’]].⁷ So if this standard view of the syntax of progressive verb phrases is correct, there are an infinitude of expressions that could fit into the relevant contexts exemplified by (2) and (3). All we need to do is replace ‘house’ with ‘big house’, ‘big red house’, ‘big red house that would look like a shopping center’, etc. In the second place, progressive constructions are also productive *over time*. When children and adults learn a new verb, they automatically know how to create the verb’s progressive form. If progressive constructions were non-productive (e.g., if they were idioms), this commonly held ability would be inexplicable. Although there are only finitely many verbs in our lexicon at any given moment, there are indefinitely many potential verbs that could enter it. And whenever one of these potential verbs actually does enter the lexicon, we automatically know how to create its progressive form. This kind of productivity shows that ‘build’ really is a constituent of ‘building’.

The third question is: even if the counterexample does undermine (1), does this really affect any of the discussions of RC? Clearly it does. Consider, for example, the passage below, taken from the literature, in which I have replaced the authors’ examples with some of our own:

[Reverse compositionality assumes that] each constituent expression contributes *the whole of* its meaning to its complex hosts. If that’s right, then if you understand ‘building a house’, it follows that you know everything you need to determine the meanings of ‘build’ and ‘house’ and ‘-ing’: in effect, the meanings of the parts of a complex expression supervene on the meaning of that expression. (Fodor and

Lepore, 2001, p. 366; ‘building a house’ is substituted for ‘dogs bark’, ‘build’, ‘house’ and ‘-ing’ for ‘dogs’ and ‘bark’).

The argument here doesn’t show that ‘building a house’ is an idiom. Clearly someone could understand ‘building a house’, and yet not possess everything needed to determine the meaning of ‘build’. As we’ve seen, a person could associate the right concept with ‘build a house’, even though she was unaware whether ‘build’ is telic or atelic. (I will have more to say about this passage in Section 3.) Thus, the present counterexample shows that a clearer formulation of RC is needed, and that not all claims about RC will remain true under such a formulation.

The fourth question concerns what the counterexample shows us. To answer this, notice how it was constructed. We found a linguistic environment that takes expressions of multiple different sorts as its input, but always produces a complex expression of a single sort as its output. In particular, the progressive can take either telic or atelic verbs, but the resulting verb phrase is always atelic. Let’s call this sort of linguistic environment ‘conflating’, because the resulting complex phrase doesn’t indicate which of the multiple possible sorts of constituents it is built out of.⁸ RC as given in (1) entails that conflating environments do not exist. But they do, as the counterexample shows. It’s worth noting that natural languages contain conflating environments of various kinds. For instance, there are conflating environments in phonology. The suffix ‘-ity’ creates abstract nouns out of some adjectives: from ‘pure’ we get ‘purity’ and from ‘nude’ we get ‘nudity’. But notice that if you know how to pronounce ‘electricity’ you don’t necessarily know how to pronounce ‘electric’ (could it be pronounced ‘electriss’?). A similar result holds for syntax. You might know that ‘Sue is proud of Jane’s friend’ is a sentence, without knowing that ‘of Jane’s friend’ is a prepositional phrase; could it be a complement clause, as in ‘Sue is proud that Jane has a friend’ or ‘Sue is proud to be Jane’s friend’?

In addition to the examples given, the counterexample shows that (1) can also be violated by finding expressions that have relatively complex meanings, only part of which is used in a given sentence. A moment’s thought suggests that there might be lots of counterexamples of this sort for particular lexical items. Assuming that there is only one word ‘way’, it is far from clear that someone who understands a sentence like ‘Jenny knows a way to solve the problem’ will thereby understand ‘way’ so fully that he also understands ‘Jenny laughed her way into the meeting’. Similarly, if there is only one word

‘each’, one could understand ‘John gave each boy an apple’ without understanding ‘each’ so fully as to also understand ‘John and Mary stood in the middle of the crowd and hugged each other’. (It might be natural to interpret ‘each other’ not as a reflexive construction, but as referring to everyone in the crowd except John and Mary; cf. ‘each of the others’.) There are, of course, subtle and difficult issues about determining when one word is semantically general and when there are in fact two homophonous words at play. But at the very least, Occam’s Razor dictates that the default hypothesis is that there is only one word ‘way’ and only one word ‘each’. But even if that’s incorrect, cases like these abound, and some of them are sure to present further violations of RC understood as (1). These sorts of examples can also be thought of as conflating environments, since the environment only uses ‘part’ of the word’s meaning, thus leaving it open whether the word has only that meaning, or whether it has a fuller, more general meaning.

The final question concerns whether the counterexample can be avoided by a fine-grained individuation of languages. This argument goes like this. Suppose a person X doesn’t think that ‘build’ is telic. Then we should regard X’s language as being slightly different from that of the (idealized) English speaker’s – according to X’s language, ‘build’ is an atelic verb. Hence, the logical forms X assigns to utterances of (2) and (3) are different from the logical forms that you and I assign to them. Moreover, X has a perfectly good understanding of (2), (3), and ‘build’ *as elements of X’s own language*. If this is right, then RC holds for X and X’s language, so the counterexample fails. This argument appears to have some plausibility. Individuating languages in this fine-grained way is a very common move in linguistics, where a language is standardly conceived as a state of the mind (or brain) of a speaker. Since X has slightly different linguistic abilities from you and I, it follows that X speaks a slightly different language. Indeed, you and I probably speak slightly different languages from one another. In this sense, ‘English’ is just a general label for some vaguely defined family of languages that we often find convenient to group together. For present purposes, though, I will use ‘English’ to refer to those languages in which ‘build’ is telic, and which has all the other properties commonly considered to be part of ‘the English language’.

There are three main problems with this argument against the counterexample. The first problem is that it appears to be empirically false. As we’ve seen, one could fully understand a complex expression like ‘building a house’ or ‘flaking the parts’ although one is

consciously in doubt about the meaning of ‘build’ or ‘flake’. Unless we interpret ‘understanding’ in some novel unintuitive sense (contrary to the argument in Section 1), a person’s language, no matter how finely individuated, may fail to be Reversely Compositional. I will develop this view more in Section 3, but for now I turn to the two other problems.

The second problem is that the argument is too powerful. The proposal is that RC holds for X and X’s language, although it doesn’t hold for X and English. But by hypothesis, X’s language is sufficiently similar to English that X can use (3) to communicate perfectly well with you and me, even though, strictly speaking, (3) for X is a different sentence (of a different language) than it is for you or me. (In contemporary linguistics, such an outcome is, in fact, not all that uncommon.⁹) But if this kind of fine-grained individuation of languages is legitimate here, then a precisely similar fine-grained individuation of conceptual apparatuses can show that Reverse Compositionality always holds for prototype theories of concepts (and stereotype theories, exemplar theories, etc.) Recall the original scenario where prototype theories of concepts seemed to lack RC: X has the concept PET FISH but does not have PET, because the latter contains some extra feature not present in the former. The strategy for preserving RC in language shows how to preserve RC in concepts. Rather than saying that X possesses PET FISH but not PET, we may say instead that X possesses neither concept. Instead, X possesses the concepts PET* and PET* FISH, the latter of which functions just like our concept of PET FISH in cognition. Thus, the fact that X *seems* to possess PET FISH but not PET does not show that X and X’s (purported) conceptual apparatus violates RC, but only that X’s conceptual apparatus is unlike yours and mine (although they bear some important similarities). So Fodor’s original argument against prototype theories of concepts doesn’t even get off the ground.

I once presented this argument to someone who argued that the example didn’t work because in the present case, the speaker would accept the inference from ‘x was building y’ to ‘x built y’, thus showing that she doesn’t really understand ‘building’. On the one hand, it’s hard to see why this should be so. After all, we may assume that the speaker associates the right concept (viz. BUILDING A HOUSE) with the target expression. We can also assume that she has a perfect grasp of this concept, as well. (Obviously, the grasp of BUILDING A HOUSE doesn’t require the grasp of English; just ask a monolingual house-builder from Guandong.)¹⁰ Worse yet, if such inferential considerations were valid in the linguistic case, they should

be valid in the conceptual case as well. But then it becomes hard to see why a thinker's *not* making the inference from PET FISH to PET wouldn't similarly show that she didn't possess PET FISH in the first place. A similar argument applies regarding other inferences that would be licensed by the background condition that the thinker could possess PET FISH but not PET. In short, if conceptual apparatuses are individuated very finely, it becomes hard to show that any theory of concepts – prototype, stereotype, etc. – is capable of violating RC. Thus RC becomes theoretically vacuous. So in order for RC to even begin to apply in the present context, conceptual apparatuses must be individuated more coarsely. So we should assume (at least for the sake of argument) that conceptual apparatuses are individuated coarsely enough that a prototype theory could declare X to possess PET FISH but not PET. But then it looks like languages will also require a correspondingly coarse individuation, which allows for the possibility that X understands 'building' but not 'build'. So the strategy of individuating languages very finely appears to fail.¹¹

There's a third reason why the strategy of appealing to finely-individuated languages fails, which I will sketch out here. The strategy relies on the following assumption:

- (5) If P and Q are sentences of distinct languages and P and Q each contain words that the other lacks, then P and Q are type-distinct entities.

Despite its *prima facie* plausibility, I think (5) is questionable, especially when we consider a fine-grained taxonomy of languages. After all, (3) as understood by X and (3) as understood by you and me are supposed to be such a P and Q, since they each contain different words both pronounced 'build'. But if the progressive construction always creates an atelic verb phrase, then the difference between the two 'build's might simply be *erased*. Of course, the generation of the two sentences would differ, since accessing different words is involved. But differing genealogies doesn't imply type distinctness – a sphere carved from a cube needn't differ from one carved from a pyramid. Similarly, the two versions of (3) needn't differ in any *relevant* sense. If a sentence is just a pair consisting of a sound and a meaning, as much current linguistic theory has it (e.g., the Minimalist program cf. Chomsky, 1995), then the two versions of (3) differ only if they have different meanings. But it's unclear why the two versions of (3) should differ at all in their meanings. In fact, it looks like their meanings should be identical.¹² If this is right, then (5) is false, so the objection to the counterexample fails. So at the very least, we can say

that a proponent of the fine-grained languages response to the counterexample must give reasons for thinking that (5) is true.

3. WHAT IS TRUE ABOUT REVERSE COMPOSITIONALITY?

In the previous section we saw that (1) cannot be a correct theory of RC. The existence of conflating environments shows that it's possible, at least in principle, for people to understand the meanings of a complex without understanding the meanings of its parts. So it's not 'true by definition that mastering a complex expression requires mastering its parts', because it's not true. But despite all this, there still *must* be something right about RC! Conflating environments may show that RC does not hold of necessity, but it's nonetheless true that very, very often we understand a complex expression only if we understand its constituents. Thus, as we look at the ordinary output of human linguistic abilities, we notice a very strong statistical generalization. Why does this generalization hold?

Before offering a solution, let us begin by getting clear on the question. We are looking for a theory of why the probability is very high that a person understands the meaning of 'triangle' given that she understands the meaning of 'This is a red triangle'. This high probability is for the most part¹³ a straightforward empirical datum that a theory of human linguistic abilities should explain. Contrast this explanandum with e.g., a Chomskyan notion of 'competence' (e.g., Chomsky, 1992). According to (one interpretation of) this latter notion, we understand sentences containing millions of words, but this understanding does not evince itself because of limitations in our 'performance' systems. This form of Chomskyan 'understanding' is not a phenomenon to be explained, but a theoretical posit made in order to produce the best overall theory of language. According to the more ordinary, more restrictive sense of understanding that we are interested in, speakers do *not* understand some of the exceedingly long and complex sentences that fall within their 'competence'. But they do still understand the meaning of the parts of the complexes they understand. Why?

I want to suggest that we can get a better grip on the present issue by looking to our language-processing mechanism instead of our grammar. Let me explain. A grammar of a language only tells us *that* expression S means M; it does not necessarily tell us *how* we arrive at that knowledge. For instance, a grammar may well contain a recursive principle such as

- (6) $\lceil P \wedge \text{'and'} \wedge Q \rceil$ is true iff P is true and Q is true.

Such a principle, even if correct from a purely linguistic perspective, says nothing about how our minds process the conjunctive sentences we perceive. If I tell you ‘The man with the big brown hat is funny and his daughter is a comedian’, you will have processed the entire first conjunct before you hear the word ‘and’, which is the word that tells you that the entire sentence will be a conjunction. Similarly, speakers typically have a hard time processing the sentence ‘The horse raced past the barn fell’ (meaning that the horse that was raced past the barn fell). But from the perspective of a grammar, this is just another sentence. It is up to a theory of the language-processor to explain speakers’ difficulties with the sentence.

I want to suggest that the fact that we typically understand a complex expression only if we understand its parts is due to some contingent facts about language processing. Although the details of how we process language is an ongoing subject of study (for reviews, cf. Townsend and Bever, 2001, ch. 4), at least this much appears to be clear. When our mind’s language processor processes a complex expression (whether in the case of speech perception or production) it accesses the various lexical items used to assemble a mental representation of the complex. RC appears to hold, I suggest, because speakers obtain a high degree of epistemic access to the lexical items used during the construction of this mental expression. Thus, RC appears to hold because in normal circumstances, one has a token occurrence of understanding a complex only if, as a part of the processing of the complex, one also has a token occurrence of understanding the complex’s parts. Our general disposition to understand a complex is thus developed at least in part by our ability to understand the complex via our understanding of its parts. Moreover, this is typically the only way we understand complexes.¹⁴

The theory just offered is in no sense an a priori truth. Our language-processing mechanisms could have been different so that RC wouldn’t even appear to hold. Here is one such scenario. Since the primary purpose of language is to communicate whole propositional meanings, the language processing mechanisms might have evolved (for reasons of computational efficiency perhaps) to provide us with epistemic access to the meanings of whole sentences only. In such a case, although we would understand whole sentences, we would have no epistemic access whatsoever to the meanings of

subsential expressions. In such a case, we would be less inclined to think that we have an introspectible awareness of our language than we currently do. The perception of a sentence might seem holistic: we would hear some noises, and then, *pow!* – a complete sentential (or at least clausal) meaning would leap into our heads. Speech production might seem a bit like walking: just as we can decide to walk across a bit of rough terrain without attending to the delicate movements of our legs, feet, and toes – and perhaps not even understanding why we do certain things to maintain balance (such as inhale and hold our breath when standing up from a sitting position) – so too in the case of language. When a person wanted to express something she would begin making sounds, with little awareness of why she was choosing those particular individual sounds, except that they were getting her to where she wanted to go, figuratively speaking. When she was done making sounds, she (and her audience) would both be aware of what she had expressed. Now, it's very likely that there are numerous good reasons why our language processors didn't develop this way: conscious access to the meanings of constituents surely makes the acquisition of new words easier, aids in on-line decisions about what exactly to communicate, etc. But the point here is that it looks like whatever truth there is behind RC is some sort of highly contingent fact about our linguistic abilities; it is not the 'true by definition' truth that Fodor suggests it is.

Interestingly, if the present view is on the right track, then RC is unlike its semantic brethren Compositionality and Systematicity. Compositionality and Systematicity both appear to be explicable by reference to a grammar of the language only. We don't need to know how we process sentences in order to confirm that complex expressions are built up out of smaller ones, or that 'John kissed Mary' is a sentence only if 'Mary kissed John' is too. However, it appears that we do need to advert to a different sort of system, the language-processing mechanism, in order to explain what is correct about RC.

Putting these ideas together, I offer the following theory of what is correct about RC: as a statistical psychological generalization that holds with great regularity, a person understands the meaning of a token occurrence of a (non-idiomatic) complex expression of her language by processing that expression's meaning out of the meanings of the constituents of the complex and their syntactic configurations. It is an empirical fact that this form of language processing enables us to understand the primitives and other constituents of the complex.

4. REVERSE COMPOSITIONALITY AS A METAPHYSICAL THESIS

In the bulk of this paper, I have considered RC as an epistemic thesis, in particular as a generalization about our understanding of complexes and their constituents. But RC has sometimes also been expressed as a metaphysical thesis. Roughly speaking, the metaphysical version of RC imposes a constraint on what it is to be part of a complex expression – the parts’ meanings must be fully determined by the complex’s meaning. This is what is intended by such claims as ‘the meanings of the parts of a complex expression supervene on the meaning of that expression’ (Fodor and Lepore, 2001, p. 366). However, this version of RC isn’t acceptable as it stands, since the counterexample applies here as well. The meaning of ‘Mary was building a house’ may be fully determined, even though it is indeterminate whether ‘build’ is telic or atelic.¹⁵ The existence of conflating environments strikes me as a serious problem for metaphysical versions of RC, insofar as they purport to give a metaphysical condition on being a complex expression. Although one may be able to salvage some statistical generalization here, it is clear that the resulting version of RC would not tell us about the fundamental *constitutive* relations between a complex and its parts. A similar problem arises for other claims Fodor and Lepore make in this vein, such as that ‘each constituent contributes *the whole of* its meaning to its complex hosts’ and ‘complex meanings don’t just supervene on but actually contain the constituent meanings’ (Fodor and Lepore, 2001, p. 366). After all, in a conflating environment like the progressive, it looks like ‘build’ fails to contribute its telicity to the meaning of the complex. In fact, it could be that the sentence-processing occurs by selecting the meaning of ‘build’ and stripping off its telic aspect before inserting the item into the mental representation of the sentence.

Another way of thinking of this last point is as follows. The claim that ‘constituents are by definition *parts* of their hosts’ (Fodor, 1998b, p. 52) appears to be unproblematically true only in a syntactic sense. A word can’t be a syntactic part of a complex and yet be somehow less than contained in the complex. However, it does appear that *syntactic* constituents need not be *semantic* constituents of the expressions in which they appear. Conflating environments provide one example, where syntactic constituents need not contribute all of their meaning to the meaning of the sentence. Expletives such as ‘it’ in ‘it is raining’ provide another, insofar as they do not contribute any meaning at all (e.g., Chomsky, 2000).

On a final note, at one point Fodor and Lepore discuss RC by saying ‘you won’t find a language that can talk about dogs barking but can’t say anything else about barking or about dogs’ (Fodor and Lepore, 2001, p. 366) Unfortunately, such a claim does not imply RC. To see this, notice that there could be languages in which ‘dog’ means *dog* in some but not all sentential environments. As long as the various sentential environments in which ‘dog’ means something other than *dog* can be recursively specified, such a language will remain Compositional (and hence, we may assume, masterable by finite beings). However, there is no reason that such a language could not contain a very large array of conflating environments for ‘dog’, with the result that one could know the meaning of many (indeed infinitely many) sentences containing ‘dog’ without knowing the meaning of ‘dog’. Such a language would not be Reverse Compositional in any metaphysical sense whatsoever. However, it would accord with Fodor and Lepore’s claim just given. Thus, while I suspect that Fodor and Lepore’s claim is a non-trivial truth about natural language, it is not a claim about RC.

5. FODOR’S ARGUMENT AGAINST PROTOTYPES

I will end by returning to Fodor’s argument against prototype theories of concepts. We saw at the outset that RC supplies a crucial premise in this argument. However, we can now see that RC cannot do the work that this argument assigns to it. As it is originally formulated, the argument works by comparing the combinatorial structure of concepts to the combinatorial structure of language. The core of the argument is that just as you can’t understand a complex expression without understanding its parts, so too you can’t understand a complex concept without understanding its parts. But we’ve seen that you can understand a complex expression without understanding its parts. Thus, the argument as presented is unsound. Now, since the argument is directed towards theories of *concepts*, not theories of language, one might hope to revive the argument by avoiding the detour into language. That is, one might maintain that RC is ‘true by definition’ of complex concepts and their constituents, even if this is not so for language. However, if one wants to defend a version of RC for thought, one will need to show that there are no conflating environments in the conceptual domain.¹⁶ Why this should be so is an utter mystery to me. It also becomes mysterious what concepts are expressed by the conflating environments of natural language.

Barring evidence to the contrary, this sort of strategy is a non-starter. More generally, it looks like the appeal to language is a crucial part of these sorts of argument. After all, we know virtually nothing of detail about the nature and structure of thoughts. So often the best we can do is *hypothesize* some interesting relation between language and thought, and explore the language in the hopes of better understanding thought. (Personally, I am skeptical of this strategy: the available evidence makes it wildly unclear what the precise relationship between language and thought is.) Thus, in the absence of further evidence it is natural to assume that if natural languages contain a certain kind of combinatorial mechanism (such as conflating environments), then thought does too. Additional considerations may serve to bar certain individual combinatorial mechanisms, but nothing seems to bar the presence of conflating environments in thought.

A final note about Fodor's use of RC is in order. Notice that the empirical phenomenon that makes RC seem so plausible is entirely epistemological. Why we typically don't understand a complex without understanding its parts is an interesting question that calls for explanation, regardless of what your favorite theories of mind and language are. But there is no corresponding purely empirical issue of why a complex's meaning 'contains' its constituents' meanings. Unless you *assume* that a complex's meaning determines the meanings of its parts, there is no explanans that requires theorizing. Moreover, this assumption is a crucial one in the debate about concepts, since advocates of stereotypes, prototypes, exemplars, recognition theories of concepts, etc. may want to deny this assumption. So in terms of the debate, RC is epistemological, not metaphysical. But Fodor's argument is about the metaphysical nature of concepts, not our epistemic access to them. So he will need to tie the metaphysical properties of these concepts to their epistemological manifestations. This won't be easy. We saw above that one needn't have epistemic access to the structure of one's sentences. But similarly, one could have epistemic access to features that are *not* part of one's sentences. Perhaps you can't (or don't) think that Mary is a mother without thinking about fathers, even if only briefly.¹⁷ Similarly, perhaps our language-processors briefly activate the word or concept for father when processing the corresponding sentence. (This sort of phenomenon is well-known from the priming studies that are often cited as evidence for prototype theories of concepts (e.g., Swinney, 1979).) In short, claims of the form 'You can't understand X without understanding Y' are independent of claims of the form 'There can't be an X without being a Y' or 'X is partly constituted by

Y' (where X and Y range over sentences or thoughts). Thus, we may safely conclude that Fodor's argument does nothing to undermine the status of any particular current theory of concepts.

6. CONCLUSION

Contrary to the assertions of Fodor and his opponents, the Reverse Compositionality of language is not a trivial thesis, but is a substantial empirical claim about the cognitive organization of language. Its most natural formulation (given in (1)) is false: we sometimes understand complex expressions without understanding their constituents. Moreover, we cannot avoid this fact by e.g., appealing to the fine-grained languages that particular individuals speak. However, Reverse Compositionality (as an epistemic thesis) does appear to hold as an empirical statistical claim about some cognitive aspects of human language processing. Such a statistical claim allows for one to sometimes understand a complex expression without understanding all its constituents. Thus, we cannot reject a theory of language understanding on the apriori grounds that it makes this outcome possible. However, any psychological theory that predicted that complexes are frequently understood in the absence of the understanding of their constituents would be empirically false. Finally, we saw that undermining a version of Reverse Compositionality like that given in (1) undermines its utility for rejecting statistically based theories of concepts.

ACKNOWLEDGEMENTS

I am grateful to an anonymous referee for this journal for a wealth of insightful linguistic and philosophical advice. This paper started its life as a commentary on a paper given by Phillip Robbins at the Pacific APA; thanks to those there for their useful feedback.

NOTES

¹ The proviso that the complex expression be non-idiomatic serves to rule out cases such as 'John kicked the bucket', which implies only that John died. One could easily understand the meaning of 'kick the bucket' without understanding the meaning of 'kick'. But since idiomatic expressions like these are not the norm, and

are only finite in number, they may safely be ignored. (Idioms such as ‘kick the bucket’ must be learned one by one, so if there were infinitely many of them, language learning would require learning and storing an infinite number of exceptions, which no finite being can do.)

- ² For work on the clarification of Compositionality and Systematicity, cf. e.g., Szabó (2000) and Niklasson and van Gelder (1994).
- ³ Nothing in the ensuing discussion depends on taking events as a crucial part of the semantics of English. The point can be made with any other framework; I speak of events only as a matter of convenience.
- ⁴ One might worry that ‘watch’ expresses a relatively unstructured type of action, in contrast to ‘build’, which is a complex sort of action. But the same point can be made with other atelic verb phrases, such as ‘play cards’ or ‘eat ice cream’ (as opposed to ‘eat the ice cream’). Cf. footnote 5 for further discussion.
- ⁵ In fact, the example in (3) suggests a related counterexample. So far, I spoken of telic and atelic verbs, but strictly speaking, it is whole verb *phrases* that are telic or atelic. The sentence ‘Mary built houses’, where the grammatical object is now an indefinite bare plural expression, does not express an endpoint to the activity of building houses. To see this, notice that we might say how long Mary’s building activity went on with a sentence like ‘It took Mary a month to build the house’, but we wouldn’t describe a similar building activity with ‘It took Mary a month to build houses’. The latter sentence can only mean that a month elapsed between some contextually specified starting point and the beginning of Mary’s building houses. So analogously to our original counterexample, it looks like a person could understand ‘Mary built houses’ without understanding ‘Mary built the house’, because the latter verb phrase was misinterpreted as being atelic. Note that none of this affects the original counterexample, since we can reformulate its point as being that ‘build’ forms a telic verb phrase when combined with a definite direct object like ‘the kitten’, whereas ‘watch’ does not. (Incidentally, if you think that ‘build’ doesn’t contain a specification of an endpoint, and that the telicity in ‘Mary built the house’ is produced in part by the non-progressive syntactic environment, then you can run the counterexample in reverse and show that one could understand ‘Mary was holding the kitten’ all the while misinterpreting ‘watch’ as a telic verb like ‘build’.)
- ⁶ One might worry about the details of Parsons semantics, arguing that it is not the verb that is telic, but rather the construction with the simple past, as opposed to the progressive. This worry misses the point. There must be something about individual verbs like ‘build’ and ‘watch’ that enable only the former to interact with the simple tenses to form only telic constructions and only the latter to form only atelic constructions. (Similar remarks apply, *mutatis mutandis*, to any of the other semantic theories that might have been used as an example.)
- ⁷ This structure presents us with a nice semantic result too, because it puts the direct object of the verb in the intensional context of the progressive, which helps to explain why someone could be building a house even though no house exists. Cf. Parsons, 1990 for an attempt to render the context created by the progressive extensional. Such a strategy, even if successful, is orthogonal to the relevant point here.
- ⁸ My notion of conflation should not be conflated with other uses of the term, which occasionally appear in the linguistics literature.

⁹ For instance, the peculiar properties of negative polarity items like ‘anymore’ are typically thought to be semantic. Virtually all English speakers find sentences like (i) to be clearly unacceptable. However, in my Midwestern dialect, (ii) is perfectly acceptable. But even so, it’s a non-trivial empirical question whether the various (classes of) English dialects assign differing interpretations to (iii)

- (i) *John dates Karen anymore.
- (ii) I’m so busy anymore that I don’t have time for coffee.
- (iii) Stan doesn’t wet the bed anymore.

¹⁰ It’s worth noting that such considerations regarding the inferential roles of concepts or expressions will not be available to an atomist like Fodor.

¹¹ Of course, this argument is non-demonstrative; needing coarseness in concept-individuation does not entail that languages will be individuated coarsely enough to legitimate the counterexample. But it places the burden of proof squarely on the opposition, and it is not at all clear how one might discharge this obligation.

¹² There are also technical reasons to suppose this. Many linguists adopt the Principle of Full Interpretation (e.g. Chomsky, 1995), according to which the syntactic structure that gets semantically interpreted cannot have any elements that are not semantically interpreted. Given Full Interpretation, it becomes plausible that one effect of the progressive construction is to simply eliminate any morphosyntactic elements that could designate telicity. If this is right, then it might be possible to prove, within the linguistic theory alone, that the two versions of (3) are identical, without appealing to subtle intuitions about meanings, concepts, and the like. (An anonymous reviewer correctly observes that from the standpoint of the Minimalist Framework, it’s hard to see what to do with a verb’s marks of telicity when the verb appears in progressive constructions. Assuming that these marks fall within the scope of a Minimalist theory, this is a problem for Minimalism generally. I suspect, although I cannot argue for it here, that the Minimalist will best account for such phenomena by appealing to either a certain degree of opacity at the lexical level, or else to some external conditions coming from the conceptual-intentional interface.)

¹³ I say “for the most part” because saying that we understand an expression only if we understand its parts presumes some means for determining what the parts of an expression are, and this sort of information comes from a linguistic theory. In addition to the semantically null elements posited by linguistic theories, there are also a number of rather fine-grained semantically contentful elements that appear in some constructions. The presence of these additional elements does not undermine my main claims here, although they do make room for further refinement and investigation of just what sorts of constituents we almost always understand when we understand complexes in which they occur. I leave such explorations for other work.

¹⁴ The case of our speaker who has no clue whether ‘build’ is telic or atelic might be an exception. In the processing of ‘Mary was building a house’, perhaps the speakers’ language system includes only a partially specified lexical entry for ‘build’, where this partial specification is enough to successfully enter into the given complex expression. There are other possibilities, too; perhaps in states of ignorance the language system simply assigns a default interpretation (either telic

or atelic) to the verb, and maintains that interpretation until it is corrected. But in either case, there is plenty of room for the speaker to understand the complex all the while being very clearly aware of not understanding the meaning of ‘build’.

- ¹⁵ This claim holds whether we are thinking of (i) an individual speaker who understands ‘Mary was building a house’, but doesn’t know whether ‘build’ is telic or atelic, or (ii) a pair of public languages, identical other than that in one of them ‘build’ is telic and in the other it is atelic. Cf. the remarks above about fine-grained individuation of languages.
- ¹⁶ If the terms are spelled out precisely, it is a theorem that if a Compositional language is not Reverse Compositional, then it contains conflating environments.
- ¹⁷ Note that it would be question-begging against some theories of concepts to stipulate that one could have the concept of a parent and not have the concept of a father.

REFERENCES

- Beaney, M. (ed.): 1997, *The Frege Reader*, Blackwell, Oxford.
- Bonomi, A.: 1997, ‘The Progressive and the Structure of Events’, *The Journal of Semantics* **14**, 173–205.
- Chomsky, N.: 2000, ‘Minimalist Inquiries: The Framework’, in R. Martin et al. (eds.), *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik*, MIT Press, Cambridge, MA, pp. 89–155.
- Chomsky, N.: 1992, ‘Language and Interpretation’, in J. Earman (ed.), *Inference, Explanation, and Other Frustrations*, University of California Press, Berkeley, pp. 99–128.
- Chomsky, N.: 1995, *The Minimalist Program*, MIT Press, Cambridge, MA.
- Dietrich, E. and A. B. Markman: 2003, ‘Discrete Thoughts: Why Cognition Must Use Discrete Representations’, *Mind and language* **18**(1), 95–119. Cf. pp. 109–110.
- Drijkoningen, F.: 1992, ‘Derivation in Syntax’, in M. Aronoff (ed.), *Morphology Now*, SUNY Press, Albany, NY, pp. 48–68.
- Fodor, J. A.: 1998a, ‘There Are No Recognitional Concepts – Not Even RED’, in J. A. Fodor (ed.), *In Critical Condition*, OUP, Oxford.
- Fodor, J. A.: 1998b, ‘There Are No Recognitional Concepts – Not Even RED, Part 2: The Plot Thickens’, in J. A. Fodor (ed.), *In Critical Condition*, OUP, Oxford.
- Fodor, J. A. and E. Lepore: 2001, ‘Why Compositionality Won’t Go Away: Reflections on Horwich’s “Deflationary” Theory’, *Ratio* **14**(4), 350–368.
- Fodor, J. A. and E. Lepore: 1998, ‘The Emptiness of the Lexicon: Reflections on James Pustejovsky’s *The Generative Lexicon*’, *Linguistic Inquiry* **29**(2), 269–288.
- Fodor, J. A. and Z. Pylyshyn: 1988, ‘Connectionism and Cognitive Architecture: A Critical Analysis’, *Cognition* **28**, 3–71.
- Frege, G.: 1892, ‘On *Sinn* and *Bedeutung*’, in Beaney 1997, 151–171.
- Higginbotham, J.: 2000, ‘On Events in Linguistic Semantics’, in J. Higginbotham, F. Pianesi, and A. C. Varzi (eds.), *Speaking of Events*, OUP, Oxford.
- Landman, F.: 1992, ‘The Progressive’, *Natural Language Semantics* **1**, 1–32.
- Niklasson, L. and T. van Gelder: 1994, ‘On Being Systematically Connectionist’, *Mind and Language* **9**, 288–302.

- Parsons, T.: 1990, *Events in the Semantics of English*, MIT Press, Cambridge, MA.
- Robbins, P.: forthcoming, 'The Myth of Reverse Compositionality', *Philosophical Studies*.
- Szabó, Z. G.: 2000, 'Compositionality as Supervenience', *Linguistics and Philosophy* **23**, 475–505.
- Swinny, D.: 1979, 'Lexical Access During Sentence Comprehension', *Journal of Verbal Learning and Verbal Behavior* **18**, 523–534.
- Tenny, C.: 1994, *Aspectual Roles and the Syntax-Semantics Interface*, Kluwer, Dordrecht.
- Townsend D. and T. Bever: 2001, *Sentence Comprehension*, MIT Press, Cambridge, MA.
- Vendler, Z.: 1967, *Linguistics in Philosophy*, Cornell University Press, Ithaca.
- Verkuyl, H. J.: 1989, 'Aspectual Classes and Aspectual Composition', *Linguistics and Philosophy* **12**, 39–94.

School of Social Science
3151 Social Science Plaza
Irvine, CA 92697-5100
Canada
E-mail: johnsonk@uci.edu

Manuscript submitted 1 April 2004
Final version received 10 January 2005