Philosophical Perspectives, 16, Language and Mind, 2002

DO ADJECTIVES CONFORM TO COMPOSITIONALITY?[1]

Marga Reimer
University of Arizona

1. Preliminaries

The other day, I informed my young daughter that I was going to *The Office Depot* to buy some printing cartridges. When I asked her if there was anything I could pick up for her, she told me that she wanted a 'red stapler'. I had no trouble understanding my daughter's request and, in fact, returned from the store with (among other things) a red stapler. Although I knew just what my daughter meant by 'red stapler', it is quite possible that prior to her request I had never actually encountered that particular phrase before. Suppose this to have been the case. How then was I able to understand what she had in mind—what she meant—by her request for a 'red stapler'? "Easy," you might think. Being a competent speaker of the English language, I have not only mastered its grammar, but I also know the meanings of tens of thousands of English words, including the words 'red' and 'stapler'. Because natural language is compositional—because (in other words) the meaning of a complex expression is a function of, and only of, the meanings of its parts (plus syntax) I was able to figure out the meaning of the adjective-noun construction 'red stapler'.[2] Because I assume that my daughter is likewise a competent speaker of the English language, one who intended to be speaking literally, I assume that she meant by 'red stapler' what the expression itself means: a stapler that is red. Hence, my ability to understand my daughter's request.

My interest in this paper will be with the simplest of adjective-noun constructions, constructions like 'red stapler', 'pretty flower', and 'round face'. The question I will be concerned with is whether or not such expressions have compositional meanings. To address this question is to address the question whether the meanings of such expressions are a function of, and only of,[3] the meanings of their constituent parts, plus the relevant syntax. I will argue that such expressions do indeed have compositional meanings, contrary to what a quick survey of the linguistic data might initially appear to show. Indeed, I will argue that there are two distinct (but related) levels of meaning associ-

ated with constructions of the sort in question, both of which are compositionally determined. Once these two levels of meaning are recognized as such, the initially troubling linguistic data fall naturally into place within a compositional framework.

The format of this paper will be as follows. After clarifying the operative notion of compositionality, I will draw attention to data that appear to undermine the presumed compositionality of natural language, and of adjective-noun constructions in particular. I will then go on to propose a two-tiered account of the meanings of such constructions that respects compositionality constraints. The account will be developed further by comparing and contrasting it with several recent attempts (compositional as well as non-compositional) to capture the meanings of adjective-noun constructions. In the penultimate section of the paper, I will attempt to show how the proposed account of adjective-noun constructions might be extended to other sorts of natural language expressions. Finally, I will conclude with a few brief methodological remarks.

2. The Principle of Compositionality

The principle of compositionality says, simply: The meaning of a complex expression is a function of, and only of, the meanings of its parts, plus syntax. Why would anyone suppose such a principle to be true of natural language? Well, everyday we produce and understand utterances of sentences never before encountered. Without the assumption that natural language is compositional, this feature of language—its so-called "productivity"—would seem rather mysterious. Suppose, for instance, that while visiting me at my home in Tucson, you ask me what I had for dinner the previous night, and I respond with an assertive utterance of:

> (1) Last night, I had four unripe pears, one cream-filled donut made by my mother-in-law, and a half stalk of steamed, organic, asparagus from the garden out back.

Doubtless, you have never heard an utterance of this particular sentence before. Nevertheless, you would presumably have no difficulty understanding my utterance. How is that possible? Only, it would seem, if you could rely on your knowledge of language—of the grammar of English and the meanings of the individual (English) words comprising the sentence uttered. Crudely put, you "assign and combine" individual word meanings so as to emerge with an utterance meaning. This is the meaning that you understand in understanding my assertive utterance of (1). Moreover, no other (non-linguistic) knowledge would appear to be necessary for an of understanding my utterance.[4] You don't, for instance, need to know whether the pears were Bartlett pears, whether the cream in the donut was flavored with vanilla, what my mother-in-law's name is, or whether I over-cooked the asparagus. Hence, the intuitive plausibility of the compositionality principle.

But what exactly are the "meanings" to which the principle refers? Different theorists have no doubt meant different, even vastly different, things by the notion of "meaning." The question here is, what kinds of complex "meanings" are arguably determined compositionally? The meanings in question must be such that, when appropriately "composed," they constitute possible objects of our understanding. For only then would compositionality have a bearing on the productivity of language.

There are at least two initially plausible candidates for compositionally determined meanings: linguistic (aka "conventional") meanings and propositional (aka "truth conditional"[5]) contents. The distinction between these two kinds of meaning—let's call it the "meaning/content" distinction—is a familiar and widely accepted one, at least among contemporary philosophers of language. It is certainly an intuitive one—one that captures the distinction between the linguistic meaning of an expression and its propositional contribution: its contribution to "what is said" by the utterance (in a context) of a sentence containing it. Indeed, the intuitiveness of the distinction is often appealed to in order to motivate the need, in semantic theory, for talk of "propositions" (in addition to sentences). Of course, like everything in philosophy, the meaning/content distinction is not uncontroversial, and even among those who accept it, there is controversy over how it is to be drawn in particular cases. The distinction is perhaps best illustrated by considering sentences containing indexical expressions, expressions the reference of which varies from one context of utterance to the next. For in cases of this sort, meaning and content seem clearly to come apart.[6] Suppose that you and I both assertively utter, in the appropriate sort of circumstances, a token of sentence type (2):

(2) I was there last night.

The sentences we utter have the same linguistic meaning, but their propositional contents are different: different things are "said." You have said *of yourself* that you were at some particular location l_1 at some particular time t_1; I have said *of myself* that I was at some particular (possibly different) location l_2 at some particular (possibly different) time t_2. The differences in what has been said can be articulated in terms of truth conditions: What you have said is true just in case *you* were at l_1 at t_1; what I have said is true just in case *I* was at l_2 at t_2. To return to sentence (1) above. This sentence contains a number of indexical expressions: 'last', 'I', 'my', 'out back'. To understand what I say in assertively uttering (1), you must arguably understand the propositional content expressed by that utterance. Such understanding involves somehow "grasping" the contents of the indexical expressions in question: the "entities" contributed by those expressions to what it is that I say. One must know who the speaker is in order to grasp the content of 'I'; one must know roughly where the speaker is in order to grasp the content of 'out back'. In order to be in a position to know such things, one must arguably understand the reference-determining rules associated with the indexical expressions. One

must, for instance, know that 'I' refers to the speaker in the context and that 'out back' refers to some proximate, though indeterminate, location outside of and behind the speaker . Nevertheless, it seems doubtful that there is any substantive sense in which, as speakers of the language (rather than as *theorists* of the language) we routinely understand, grasp, or "entertain" compositionally determined sentence types, where these are thought of as being composed out of (*inter alia*) reference-determining rules of the sort associated with indexical expressions like those just considered. After all, outside academia, how often do people ponder the "meaning" of sentences types?[7] Hardly ever. I am not denying that speakers (in some sense) "grasp" the linguistic meanings of the individual expressions comprising the uttered sentences. I am denying that, *in addition to* grasping these linguistic meanings, speakers *also* grasp their concatenation—the sentence types themselves.[8] The hypothesis sentence types are grasped in understanding particular utterances of sentences contributes nothing to an explanation of the understanding of the utterances themselves. Moreover, "productivity"—however construed—is of course concerned with *products* of some sort—and utterances (in contrast to sentence types) are surely products; they are produced by speakers and consumed by hearers. Because propositional content appears to be relevant to the understanding of *utterances* (and thus to productivity), it is this notion of "meaning" with which I will be primarily concerned. That is, I will be focusing on the putative compositionality of propositional content, not of sentence-type meaning. And the question, for the moment, is the following: Is the content of an adjective-noun construction—its contribution to the proposition expressed, to "what is said"—compositionally determined? Are such contents a function of, and only of, the contents of their constituent parts (plus syntax)?

3. The Linguistic Data

If we make the mistake of hastily generalizing from the 'red stapler' example described above, we might well be led to the false conclusion that the contents of adjective-noun constructions are obviously compositionally determined. I know what 'red stapler' means and, in order to know this, I need only know that 'red' means red, that 'stapler' means stapler, and that the former expression modifies the latter. However, a brief survey of the relevant data indicates that such a conclusion would be much too quick. Even if, in the end, a case can be made for the compositionally of all adjective-noun constructions, it will not be an easy case to make. Because the apparent non-compositionality of adjective-noun constructions is well-known, I will be brief in my presentation of the data supporting it, drawing heavily on Lahav's (1989) paper on the topic.

Consider assertive utterances of the following ten sentences, and ask yourself the following question: What is the "content," the "propositional contribution," of the adjective-noun construction, as that expression occurs in the utterances in question? Less technically, what do the constructions in question contribute to "what is said" by those utterances? Imagine the contexts of utter-

ance to be the most ordinary of such contexts, contexts in which you might actually encounter utterances of the following the sentences.

(3) I prefer the taste of *red apples* to the taste of green apples.
(4) *Red watermelon* is so refreshing.
(5) Fred wore a really *nice tie* to his wedding.
(6) My brother is a *nice person*.
(7) I never really was a *good student*.
(8) I like to think that I'm a pretty *good person* though.
(9) I bought a real *slow oven* from GE.
(10) After a glass or two of wine, I become a *slow thinker*.
(11) I have a *strong friendship* with only a few people.
(12) After all that walking to and from campus, I have *strong legs*.

Prima facie, the facts seem to be (roughly) as follows: the content of 'red apples' is *apples that are red-skinned*; that of 'red watermelon' is *watermelon that has red pulp*; the content of 'nice tie' is *tie that is attractive*; that of 'nice person' is *person who is kind and thoughtful*; the content of 'good student' is *student who does well in his/her courses*; that of 'good person' is *person who 'does the right thing';* the content of 'slow oven' is *oven that takes noticeably longer than most reach the set temperature*; that of 'slow thinker' is *thinker who takes noticeably longer than most to figure things out*; the content of 'strong friendship' is *friendship that is likely to endure*; that of 'strong legs' is *legs that have significant muscle mass*.

What do these putative facts suggest? They suggest that the content of adjective-noun phrases is not (at least not always) compositionally determined. In other words, the propositional content of such phrases appears to somehow go beyond the contents of the individual (syntactically combined) components. At least in some instances. Thus, while the content of 'red apples' appears to be *apples that are red-skinned*, nothing in the concept of 'red' or in that of 'apple' suggests that the color of an apple is identified with the color of its skin. And while the content of 'good student' appears to be *student who does well in his/her courses*, nothing in the concept of 'good' or in that of 'student' suggests that the quality of a student is determined by the quality of his/her course performance.

So much for the linguistic data, now let's have a closer look at the problem, the "puzzle," allegedly posed by that data.

4. The "Problem" of Adjectives

As with all philosophical problems, this particular one involves a puzzle of sorts. In understanding utterances of sentences containing adjective-noun constructions, we somehow figure out what those constructions mean—what they contribute to what the utterance itself means. Or so it appears. Now if the contents of such constructions were determined compositionally, then our ability

to figure out those contents would be easily enough explained: we would simply calculate and syntactically combine the contents of the individual components. But it appears that—in some cases at least—the contents of adjective-noun constructions are not determined in this manner. Sometimes, the whole is greater than the (syntactically combined) parts. Or so it appears. For instance, I know (somehow) that a 'good student' tends to do well in his/her courses—even if his intentions are questionable. And I know that a 'good person' tends to have the right sorts of intentions—even if he doesn't ever do well in his courses. How do I figure out the contents of such expressions, if such contents are not compositionally determined?

5. Preliminary Sketch of A Solution: Two Kinds of Meaning

In order to solve the problem at hand, we need to return to the distinction drawn above between linguistic meaning and content. Let me refer to these two kinds of meaning as "L(linguistic)-meaning" and "C(content)-meaning." Roughly, L-meaning is *conventional meaning*, where the conventional meaning of an expression is the sort of meaning that a lexicographer compiling a dictionary attempts to capture, one based on standard use within a given linguistic community. C-meaning (again, roughly) is *content*,[9] where the content of an expression is its contribution to the proposition expressed (to what is said) by an utterance of the sentence in which it occurs. What I have in mind here corresponds at least roughly to what is often called "semantic value." Complex meanings, whether linguistic or propositional, are concatenated in accordance with relevant syntax/logic. Meaning and content are related to one another by *context*: in ideal circumstances, the hearer is able to discern an expression's content on the basis of his knowledge of linguistic meaning and contextual facts. As I intend the notion of "context," it includes everything relevant to the interpretation of the utterance. Following Szabo (2001) I construe context as,

> a wide and heterogeneous collection of facts concerning the linguistic and non-linguistic environment of a particular use of an expression. It includes facts about the time and the location of the utterance, facts about the speaker, the hearer, and the salient objects around them, facts about their shared background knowledge, about the form and content of the conversation they had before the utterance was made, and perhaps much more… (Szabo, p. 120)

It would of course be possible to say a great deal more about all three of these notions—linguistic meaning, propositional content, and context. But I want to say only as much as necessary in order to lay the ground-work for the proposed account of adjective-noun constructions. In this way, I will be able to remain as neutral as possible with respect to precisely what each of these semantic notions involves.

Now let's attempt to apply the meaning/content distinction to simple adjective-noun constructions of the sort in question. Here's a preliminary pro-

posal. Generally speaking, the L-meaning of an adjective-noun construction *AN* is compositional: it is simply *N which is A*.[10] The C-meaning of such a construction is *BN*, where adjective *B* specifies the *contextually sensitive* conditions, both necessary and sufficient,[11] for the applicability of *A* to *N*. In most—perhaps in all—cases there may be several different ways of characterizing *B*. This is all very abstract. Let's apply it to a few examples. Let's assume that the context is the most ordinary of contexts. Suppose that the speaker and hearer are discussing their preferences with regard to apples. Now the color of an apple is ordinarily identified by the natural color if its skin. Since our context is (by hypothesis) an ordinary one, let's assume that (in this context) the color of an apple is indeed identified by the natural color of its skin, and that both speaker and hearer are aware of this fact. Then, the L-meaning of 'red apples' is *apples that are red*, the C-meaning of 'red apples'—as it occurs in an assertive utterance of (3)—is *apples that have (naturally) red skins*. Alternatively, the C-meaning might be characterized as: *red-skinned apples*, *apples that are red on the outside*, *apples with red exteriors*, etc. The L-meaning tells you (in effect) to apply the term to apples that are red, and (in a context of the sort in question) these happen to be those with red skins. Now let's suppose again the context to be an ordinary one, one in which the speaker and hearer are discussing their pasts as graduate students. Then, the L-meaning of 'good student' is *student who is good*; the C-meaning of 'good student'—as that expression occurs in an assertive utterance of (7)—is *student who does well in his/her courses*. The L-meaning tells you to apply the expression to students who are good—and, as both the speaker and hearer know, these are (in the context in question) one and the same as students who do well in their courses.

Now let's change the context to a extraordinary (counterfactual) one in which all apples are red-skinned, with pulps that are either white or red. Suppose once again that speaker and hearer are discussing their preferences with regard to apples. Then, the L-meaning of 'red apples' will be the same as in the previous case—*apples that are red*, but the C-meaning of that construction—as it occurs in an assertive utterance of (3)—will more than likely be *apples with red pulp (with red interiors, insides*, etc.). After all, in such a context, talk about apples would surely be facilitated were the color of an apple identified by the color of its pulp, rather than by the color of its skin. Now suppose that in our hypothetical context students do not take courses, but instead attend intermittently held lectures given by the leading intellectuals of the day. Suppose further that these students are judged by others solely on the basis of the quality of the questions they ask during lectures. Then, the L-meaning of 'good student' will still be *student who is good*, but the C-meaning of that construction—as it occurs in an assertive utterance of a sentence like (7)—would probably be something like *student who asks thoughtful/intelligent questions*.

How does this generate a solution to our problem: the problem of reconciling the apparent non-compositionality of adjective-noun constructions with the undeniable facts of productivity? Competent speakers of the language know the L-meanings of *AN* constructions; this is just part of what it means to be

linguistically competent. Such speakers are also generally aware of the contextually sensitive conditions under which adjective *A* applies to an *N*.[12] If they are not aware of these conditions prior to the utterance, they can often figure them out on the basis of their knowledge of the relevant empirical facts. Thus, suppose that I have never heard of a 'blue mango', and my sister tells me that she recently tried one. I know (in virtue of my linguistic competence) that 'blue mango' refers to mangos that are blue. I might do a bit of analogical reasoning to determine that what my sister has in effect said is that she recently tried some mango that had *blue pulp*. After all, I know that a mango is a kind of fruit and I know that the color of a fruit is sometimes identified with the color of its pulp (as in 'red watermelon'). Now of course, I might be wrong—perhaps the color of a mango is identified by the color of its *skin* (as in 'red apple'). But this doesn't affect the point that I am making, which is that hearers arrive at an interpretation of the content of adjective-noun constructions by employing their knowledge of the construction's linguistic meaning, along with their knowledge of—or beliefs about—the contextually sensitive conditions under which *A* applies to an *N*. The compositionally calculated C-meanings are incorporated into the overall (literal[13]) utterance meaning. Assuming there is no difficulty understanding the rest of the utterance, the utterance as a whole is understood.

I would now like to develop the proposed view a bit more and do so by comparing and contrasting it with other recent accounts of the linguistic behavior of adjective-noun constructions—accounts developed in light of the data discussed in section 3.

6. Working out the Details: Comparison with Other Accounts

(a) Lahav

Over a dozen years ago (in 1989), Ran Lahav published a paper in which he argued that linguistic data of the sort we've been looking at threatened compositionality. In particular, he claimed that the data strongly suggest that the propositional content of an adjective varies with the noun it modifies. Thus, in some linguistic contexts, such as 'red apple', the content of 'red' is *red on the outside*; in others, such as 'red watermelon', it is *red on the inside*. I agree completely with these observations. What I don't agree with is Lahav's claim that such observations undermine the view that adjective-noun constructions conform to compositionality. More specifically, I disagree with his contention that the data undermine any substantive compositionality thesis; that is, any compositionality thesis of the sort designed to explain productivity.

I think that Lahav is led to error because he fails to distinguish between L-meaning and C-meaning. He correctly notes that the content of an adjective (which he defines in terms of its "applicability conditions") varies from one linguistic context to another. Lahav's mistake is in assuming that *content* must be uniform across contexts, if adjectives are to conform to compositionality (in any substantive sense). But this is a mistake. *L-meaning* (or conventional mean-

ing) is what is uniform across contexts; content varies across contexts. This is just as it is with run-of-the-mill indexicals like 'I': L-meaning is constant. That's what it means for the expression to be *univocal* (rather than ambiguous). In contrast, the C-meaning of an indexical expression varies with context. Such is the nature of an indexical (context sensitive) expression—whether it be an adjective, a pronoun or an incomplete definite description.[14]

That both types of meaning, conventional as well as propositional, are compositional seems plausible. *AN* L-means *N which is A* and in all instantiations of this schema, *A* L-means the same thing. What does it mean? How could it be other than the obvious: *A*? And what does *A* mean? Whatever exactly *A* means, a speaker's understanding of that meaning is obviously going to be derived (somehow) from his/her previous encounters with applications of *A*. This suggests that the L-meaning of *A* is plausibly viewed as an *abstraction* from community-wide applications of *A*. (This accords with the construal of L-meaning as the domain of the professional lexicographer.) Now for a few illustrations. 'Good student' L-means *student who is good*, 'good person' L-means *person who is good*, and in both cases 'good' has the same meaning: *good*. To determine the C-meaning of the construction, you combine the C-meanings of the parts. In the case of 'red apples', the C-meaning, in the first context considered above, is *apples that are red-skinned*. The content of 'red' is *red-skinned* the content of 'apples' is *apples*. In the case of 'good student', the C-meaning, in the first of the two contexts considered above, is *student who does well in his/her courses*. The content of 'good' (in that particular context) is *the property of doing well in one's courses*, the content of 'student' is *student*. Assign and combine these contents and you emerge with the content of the whole.

On behalf of an approach like Lahav's, one might make the following response. What's really interesting about compositionality is its connection to productivity. The central question here is the following: Is linguistic understanding a mechanical (compositional) process or a heuristic one? The test cases here are those involving our understanding of utterances of *novel* sentences. And if we reflect on how we arrive at such meanings, it seems clear that the process is not a mechanical one—it is not an "assign and combine" process of the sort associated with compositionality constraints. As Lahav plausibly observes:

> Psychologically speaking, it seems quite clear what the cognizer does when applying an adjective in a newly encountered linguistic context: he uses analogies or similarity relationships to go from familiar linguistic contexts to new ones. One can figure out that, for example, a red box is red on the outside by making an analogy—and not a trivial one—to red houses and red balls, but not to red crystals or watermelons. (Lahav, p. 278)

I agree with all of this—or at least I am not committed to denying any of it. For it does nothing to undermine an "assign and combine" approach to propositional content—a compositional approach. All that Lahav's reflective obser-

vations suggest is speakers sometimes appeal to analogies in order to figure out the content, the C-meaning, of an adjective. To figure out the content of 'red' in as it occurs in 'red box' (as uttered in a context where boxes are generally closed, and vary in external color), it might help to be aware of the relevant sorts of similarities to red balls and red houses. In all three cases, the color of *N* is likely to be identified with its external color. The reason has to do, naturally enough, with communicative efficiency. Often, one simply does not have access to the internal color of houses, balls, or boxes. Moreover, the external color of such things (which typically is accessible to speakers) varies widely among things of the same type. Such variation allows for a convenient mode of discrimination among houses, balls, and boxes. (It is true, of course, that we often don't have access to the interior color of a grapefruit, but since all grapefruits have the same external color, the expression 'yellowish grapefruit' would be about as useful as 'striped tiger'.) But again, none of this is at odds with the view that contents are assigned and combined in a fairly mechanical manner. What's heuristic, if anything is, is the means by which we determine the content of an adjective, as it occurs in the use of a particular adjective-noun construction.

(b) Fodor

In a recent (1998) paper, Fodor attempts to salvage compositionality in light of the troubling data. He basically claims that the meaning of *AN* is simply *N that is A for Ns*. This is true for any instantiation, and in all instantiations *A* means the same thing: *A*. My central concern with Fodor's account is his apparent assumption that compositionality of L-meanings explains productivity. In particular, I take issue with his assumption that the syntactically concatenated linguistic meanings are what we produce and understand—that they are the objects of "productivity." To see the problem with Fodor's assumption, let's return to an utterance of 'I was there last night'. First, suppose that you know who is speaking, you know the place referred to, and know (in some sense) the day of the utterance. Then, you understand the utterance, if you are a competent speaker of the language. That's obvious, if anything is. Now suppose you have no idea who the speaker is, no idea where the specified location is, and no idea when the sentence was uttered. Perhaps the sentence (token) is inscribed on a bathroom stall. Now I suppose there is a sense in which you nevertheless "understand" what the sentence says: (something like) the speaker was at the specified location the night before last. Indeed, if you are an analytic sort, you might even be able to spell this out (as I have just done). But to understand this concatenation of linguistic meanings, is not to understand anything other than a theorist's artifact: a concatenation of linguistic meanings not embedded in any particular (non-linguistic) context. Thus, it seems to me that Fodor effectively conflates propositional content with sentence meaning. For it is the former, and not the latter, that constitutes the objects of productivity.

To see the intuitive implausibility of Fodor's conflation, consider an utterance of the following sentence (as produced by the author):

(13) My husband is a large white man with red hair and a square jaw.

You presumably understand *something* when I utter this sentence. What exactly is it that you understand? Is it anything remotely like:

(14) My husband is large *for a man*, white *for a person*, has hair that is red *for hair* and a jaw that is square *for a jaw*.

Suppose that this is indeed what you understand—and all that you understand. (Perhaps through some bizarre turn of events you have never seen people before.) In that case, the following is possible on Fodor's account. Although you understand what I *say* (specified by (14)), you have virtually no idea of what I *mean*, of what I intend to communicate—even though I am speaking *literally*! You might (for example) speculate that I mean that my husband is well over 50 feet tall, that his entire body is white as snow—inside and out, that his hair is fire engine red—inside the shaft, that his jaw is such that a sketch of it could provide a geometer's illustration of a (perfect) square. Any view that has this as a consequence should be rejected.[15] As Strawson (1950) taught us in "On Referring," explicating the context-independent meaning of a sentence—as (14) arguably does—is not specifying what is asserted, the latter of which is intrinsically context-based.

An advocate of an account like Fodor's might respond as follows. Language is essentially about things linguistic; once we start talking about such contextually-enriched entities as "contents" and "propositions," we are talking about something other than language proper. So, what's key in talking about language is L-meaning—which is compositionally determined. Moreover, to assume the compositionality of L-meaning is necessary if we are to explain productivity—even if we construe the objects of productivity as contextually-infused.

I have no quibble with any of this. The point is that compositionality of L-meanings does not—by itself—explain productivity. After all, compositionally determined L-meanings are not really products—they are neither produced nor "consumed" by the ordinary speaker. (Perhaps they produced and consumed by certain theorists, though.) At any rate, in order to explain the productivity associated with utterances of sentences, we need to bring in empirical non-linguistic knowledge. Moreover, by bringing in contextual considerations of the sort in question, we are able to offer a natural explanation of a fairly common phenomenon that might seem puzzling to a theorist who focuses on L-meaning to the exclusion of C-meaning. For instance, suppose that while shopping with my five year-old daughter, I ask her to pick out a 'nice red watermelon'. It is not inconceivable that she might respond with an exasperated utterance

of "There are only green ones in this store."[16] Suppose that she does just this—despite the mounds of red watermelon before her. Well, we want to say that she has made a mistake of some sort. The question is: Is it a *linguistic* one? Doesn't she know what 'watermelon' means? Of course she does—she has no problem identifying the large 'green' melons in the store as watermelons . Doesn't she know what 'red' means? Of course she does—as she'll tell you, it's her favorite color. Asked for examples of things that are red, she will readily provide you with a list such as the following: fire engines, tomatoes, Santa's suit, and her new stapler. What, then, is the problem? She knows the meanings of 'red' and 'watermelon', and yet she can't recognize a red watermelon when she's looking right at one. The problem is that she doesn't understand my utterance because she is not in a position to calculate the C-meaning of 'red' as it occurs in my utterance of 'red watermelon'. She knows its L-meaning, but because she does not know that the color of a watermelon is ordinarily identified by the color of its pulp, she cannot see that red watermelons are all around her. (And this doesn't mean she's merely being dim; there's certainly no obvious reason why the color of a watermelon should be identified in this way. After all, all watermelons have red, or pale red, pulp. The reasons for the convention can be figured out though: watermelon is typically displayed with its bright red interior fully exposed and with its pale green skin barely noticeable. Moreover, the quality of a watermelon is more likely indicated by the color of its pulp than by the color of its skin.)

Fodor could of course attempt to explain the data without the terminology of the proposed view. He could simply claim that the fact that my daughter understands the meaning of 'red watermelon' is no guarantee that she will recognize one when she sees one. My objection to this way of putting things is that it suggests falsely that what ordinary folk (like my daughter) understand are L-meanings, rather than C-meanings. But it is surely the latter than constitute the "products" of productivity—the "items" that are produced and consumed by ordinary speakers of the language. It also suggests that context has nothing to do with how we understand utterances—which is surely, in some deep and important sense, just plain false.

(c) Szabo

Szabo (2001) formulates the compositionality principle in terms of what he calls the "Context Thesis," and then goes on to develop an account of adjectives that is consistent with it. According to the Context Thesis, the content of an expression depends on context only insofar as the contents of its constituents do. Szabo proposes an account of adjectives according to which most such expressions are "incomplete," containing a hidden variable in their logical form, to be filled by a contextually specified value. Thus to say that an apple is red is to say that there is some contextually specifiable part of it that is red. What is that part? In ordinary cases, the outer surface. To say that a person is

good is to say that there is some contextually specifiable aspect of that person that is good. What aspect is that? Generally, at least, the aspect in question would be the person's moral character.

There are a number of advantages to Szabo's account. In particular, it captures intuitions about both context-sensitivity and stability of meaning. This is clear from its treatment of utterances of sentences like (3)–(12). Although the word 'red' seems to contribute something constant in all of its occurrences (the property of being red), it also seems to contribute something more—something that varies with the noun it modifies (the respect in which the red object is red). It also captures truth intuitions. Assertive utterances of sentences like (3)–(12) emerge with the truth values they intuitively appear to have. Moreover, the account preserves compositionality, as captured by the Context Principle. In this way, we can retain the currently popular and not implausible view that the "proposition expressed" by the assertive utterance of a declarative sentence is compositionally determined: determined by the syntactic combination of the semantic values of the individual components. We are thus able to provide a systematic account of the proposition expressed by utterances of sentences containing constructions of the sort in question. And this in turn would appear to allow for a plausible account of the productivity of language—the ability of speakers to produce and understand utterances of novel sentences.

The proposed account has these virtues as well, and so the question becomes: How is one to choose between the two accounts? There are certainly important differences between the over-all approaches of the two accounts. For one, Szabo accommodates the context-dependence of adjectives by the positing of a hidden variable, whereas I do not. Instead, I view the context-dependence of adjectives as just one instance of what Kent Bach (1994) has referred to as "semantic underdetermination" at the lexical level. The presence, in an utterance, of a semantically undetermined lexical item simply means that "pragmatic specialization" is needed for a determinate proposition to be expressed. Thus, what is required in interpreting utterances containing adjective-noun constructions like 'red apple' or 'nice person' is not disambiguation of the adjectives, but their pragmatic specializations. On the proposed view, pragmatic specialization is effected by appeal to the context, broadly construed. Thus, from the point of view of the proposed account, the postulation of hidden variable is superfluous. Indeed, such postulation effectively ascribes to adjectives a logical complexity that they simply do not have. This relates to a second important difference between the two accounts: Szabo sees the context-dependence of adjectives as not widely generalizable, whereas I see such context-dependence as just one instance of a pervasive feature of natural language. (More on this below.) A third difference is that Szabo's account is silent on the relation between meaning and content, whereas the proposed account is not. By means of this distinction, I am able to unify the data associated with utterances of sentences like (3) through (12). In all cases, such utterances are understood by means of (i) knowledge of semantically underdetermined lexical meaning and (ii) knowl-

edge of a wide range of contextual factors. The former is acquired via previous encounters with applications of the expression in question; the latter is simply assumed (non-linguistic) background knowledge. By means of (i) and (ii) the speaker is able to discern the content of the expression, as it occurs in a given utterance. How this process of discernment occurs is an interesting issue; but it is the subject matter of pragmatics and psycholinguistics, rather than of philosophy of language proper.

7. Other Expressions

This paper has focused exclusively on the simplest of adjective-noun constructions and the apparent problem that they pose for compositionality. A natural question at this juncture is whether or not the proposed compositional treatment of such expressions can be extended to other sorts of expressions. I think that the answer is 'yes'. More cautiously, to the extent that other sorts of expressions pose, or appear to pose, similar problems for compositionality, the proposed treatment of adjective-noun constructions can be extended to them.

Let's look at some examples involving verbs to see how this might be done. Consider the highlighted expressions as they occur in each following six sentences.

(15) The washing machine *ruined* the sweater I bought.
(16) You *ruined* the plans for Mark's surprise party.
(17) Would you mind *giving* me a glass of water?
(18) That remark just *gave* me a great idea!
(19) Fred *opened* the door for his mother.
(20) I *opened* my eyes.

Lahav argues that examples such as these show that verbs behave like adjectives—in so far as their applicability conditions vary with the linguistic context. To ruin a sweater is to damage its fabric; to ruin a surprise party is to let the word get out; to give someone a glass of water is to hand to them; to give someone an idea is to say something that causes them to entertain that idea; to open a door is to is move it so as to enable entry into the room that it conceals when closed; to open one's eyes is to move one's eyelids in an upward direction. The point should be clear—the applicability conditions of verbs do indeed vary with linguistic context. But this does not mean that they do not behave compositionality. More specifically, it does not mean that they do not behave compositionally in a substantive sense that sheds light on productivity. The propositional contribution of a verb varies with linguistic context, yet the linguistic meaning of that verb does not. The hearer is able to interpret utterances of sentences like (15)–(20) by conjoining his knowledge of the (semantically indeterminate) lexical meaning of the verb with his knowledge of (or beliefs about) the conditions given which a given verb applies to a given activity. The former is derived somehow from previous encounters involving appli-

cations of the verb; the latter is either known from past similar applications or discerned via some sort of analogical reasoning from past applications of the verb.

This quick look at the context-sensitivity of verbs is intended only to suggest that the sort of context-sensitivity that infects adjectives may well be a pervasive one.

8. Concluding Remarks

Philosophers of language are notorious for making distinctions. Over the years, they have posited distinctions between: sense and reference, speaker meaning and expression meaning, speaker reference and semantic reference, what is said and what is implicated, wide and narrow scope, *de dicto* and *de re* readings, and character and content. One could go on and on. These distinctions are posited for reasons—there are data that need to be explained and the posited distinction is thought to do the work. What is interesting is that, in attempting to accommodate the data associated with productivity, distinctions of the sort in question are not (so far as I am aware) generally invoked. But as I hope to have shown in this paper, the puzzling data concerning adjectives is easily accommodated within a compositional framework (indeed, within a two-tiered compositional framework)—provided we remain sensitive to the distinction between what I have called "L-meaning" and "C-meaning."

Notes

1. I would like to thank Kent Bach, Dave Chalmers, David Copp, Ray Elugardo, Ernie Lepore, Mark Siebel, Jason Stanely, and Zoltan Szabo for helpful comments on earlier drafts of this paper. I would also like to thank participants at Chris Gauker's Context Conference, held at the University of Cincinnati in November of 2001, where this paper was originally presented.
2. For a classic example of this sort of argument for the compositionality of natural language, see Davidson (1965); for a more contemporary example, see Fodor (1998).
3. Presumably, no one doubts that component expression meanings and syntax *contribute* to the meaning of a complex expression; the question is whether they are the *sole* contributors to such meaning.
4. This is not quite accurate, as will become evident in the subsequent discussion of indexical expressions.
5. Although content is often thought of in terms of *truth conditional* content, my use of the term is broader, allowing for talk of the content of expressions in uses of non-declarative sentences.
6. I do not intend to suggest that meaning and content always, or even generally, come apart; only that they sometimes do. Occasional factual divergence is enough to establish theoretical distinctness.
7. The academic ponders the meaning of a sentence-type when he asks such questions as: What does such-and-such a sentence mean—*independently of* any particular context of utterance?

8. To make this move without argument would be to commit (something like) the fallacy of composition.
9. The "C" can also be thought of as representing *context*-sensitivity.
10. There are exceptions to this generalization. For instance, 'coffee cup' does not mean *cup that is coffee*. But I think that at least some of these cases can be dealt with simply by turning the adjective into an appropriate (grammatical) construction, as in *cup that is coffe-ish*.
11. I specify the conditions as necessary and sufficient in order to accommodate the apparent determinacy of propositional content. To the extent that such content might turn out to be indeterminate, the necessity and sufficiency of the conditions can be appropriately relaxed.
12. Is this also part of linguistic competence? I am not sure what to say here, but think the question may be largely verbal.
13. The proviso here is meant to exclude from the meaning of an utterance such things as conversational implicatures.
14. I don't mean to suggest that pronouns and adjectives are "indexical" in any sense other than this.
15. Mark Siebel (2001) makes the same basic point against Fodor's view.
16. Conceivable, but unlikely. She actually told me that watermelons were "green and red."

References

Bach, K. 1994: Conversational Impliciture. *Mind and Language* 9, 124–162.

Davidson, D. 1965: Theories of Meaning and Learnable Languages. Reprinted in *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press, 3–15.

Fodor, J. 1998: There are no recognition concepts; not even RED. Online paper, http://www.nyu.edu/gsas/dept/philo/courses/concepts/Fodorch4.htm

Lahav, R. 1989: Against Compositionality. The Case of Adjectives. *Philosophical Studies* 57, 261–279.

Siebel, M. 2001: Red Watermelons and Large Elephants: A Case Against Compositionality? *Theoria*, vol. 15/2, no. 38.

Strawson, P. 1950: On Referring. *Mind*, 320–344.

Szabo, S. 2001: Adjectives in Context. In *Perspectives on Semantics, Pragmatics and Discourse: A Festschrift for Ferenc Kiefer*, R.M. Harnish and I. Kenesei, eds.) Amsterdam: John Benjamins.