

RED WATERMELONS AND LARGE ELEPHANTS: A CASE AGAINST COMPOSITIONALITY?

Mark SIEBEL*

* Research Group *Communication & Understanding*, University of Leipzig, Burgstr. 21,
04109 Leipzig, Germany. E-mail: siebel@rz.uni-leipzig.de

BIBLID [0495-4548 (2000) 15: 38; p. 263-280]

ABSTRACT: The standard argument against the compositionality of adjective-noun compounds containing "red" says that "red" does not make the same semantic contribution because a red car has to be red outside whereas a red watermelon has to be red inside. Fodor's reply to that argument is that the inside/outside feature is semantically irrelevant because "red *F*" just means *F which is red for Fs*. That account agrees with our intuitions concerning analyticity; but it seems to be in conflict with a central test for understanding: a person who knows nothing else about these expressions than what is offered by Fodor is far from applying them successfully.

Keywords: adjective-noun compounds, compositionality, linguistic competence, meaning, semantics.

CONTENTS

Introduction

1. The Principle of Compositionality
2. The Standard Argument against Compositionality
3. Fodor's Account
4. The Pros and Cons of Fodor's Account
5. Conclusion

Bibliography

Introduction

In this contribution, I examine whether the semantics of certain adjective-noun compounds threatens the principle of compositionality in its application to natural languages. Initially, I confine my attention to compounds containing the adjective "red", such as "red car" and "red watermelon". Later on, I will also include the expressions "large elephant" and "large ant" because they have quite a bit in common with the former.

In the first section, I present the basic idea of compositionality, one of the main tasks it is supposed to fulfill and some of its limitations. Fur

thermore, I try to show why it is plausible to assume that the meaning of adjective-noun compounds containing "red" is in fact compositionally determined.

In the second part, I offer the standard argument against a compositional semantics of these expressions. It claims, roughly, that the adjective "red" does not make the same semantic contribution to the compounds in which it is embedded because what counts as red in each case is different from noun to noun. A red car, for instance, has to be red outside, whereas a red watermelon has to be red inside.

In the third section, I present Jerry Fodor's reply to that argument. It is rather radical because it does not accept that the inside/outside feature is part of the meaning of "red car" and "red watermelon". According to that reply, these compounds are on a par with "large elephant" and "large ant". "Large elephant" just means *elephant which is large for elephants*; and the same holds for "red watermelon": it means nothing else than *watermelon which is red for watermelons*. Hence, there is no threat for compositionality because the inside/outside aspect is semantically irrelevant.

In part 4, I point to the pros and cons of that account. There are good arguments in favour of it; but there is also a strong argument against it. A person who knows nothing else about these compounds than what is required by Fodor's proposal will not pass a very common and central test for understanding: such a person is far from applying these expressions successfully.

In the end, section 5, I cannot help but leave open the answer to the question in the title. I think, however, that the way to that undecided conclusion reveals some important insights which are helpful to further reflections on compositionality.

1. The Principle of Compositionality

In its intuitive form, the principle of compositionality claims that the meaning of a complex expression is a product of the meanings of its constituents and their mode of combination. In other words, the meaning of a complex expression depends on nothing else than the meanings of the expressions contained in it and the way in which they are combined.¹

As John Lyons says, the principle of compositionality "is absolutely central in modern formal semantics".² Moreover, it seems to be indispensable in a treatment of *natural* language because it is the best, and perhaps the only, explanation available for what Jerry Fodor and other people call

productivity.³ We are able to understand an indefinite amount of novel expressions. There is, in other words, a huge class of complex expressions we can understand although we never before encountered them. This, however, seems to be possible only if our language is compositional. Compositionality is, at least, a very good explanation for that phenomenon. We can know the meaning of novel compounds because their meaning depends just on the meaning of their constituents and their mode of combination. Since we are familiar with the meaning of the parts and the grammatical structure of the whole expression, we can figure out its meaning although we hear or read it for the first time.

Of course, the principle of compositionality has to be restricted because there are complex expressions with a meaning which is clearly not a function of the meaning of their parts. One example are expressions with an idiomatic meaning, such as the phrase "kick the bucket". The idiomatic meaning of that phrase neither contains the meaning of "kick" nor the meaning of "bucket". Other examples are provided by many noun-noun compounds, such as "fireman". "Fireman" means something like *man who fights fire*. Its meaning contains, besides the meaning of "fire" and "man", an additional element, namely the semantic content of the verb "fight". That element is neither contributed by the meaning of its constituents nor by their syntactic structure.

Nevertheless, the principle of compositionality seems to have a wide field of application. There are, for example, many cases of adjective-noun combinations where it is extremely plausible that their meaning is compositionally determined. To have a concrete example at hand, consider the expression "red car". A red car is nothing else than a car which is red. Something is a red car if and only if it is red and a car. In other words, "red car" seems to be a so-called *intersective* adjective-noun compound: the extension of that compound is given by the intersection of the adjective's and the noun's extension. And this is a strong indication to the fact that the meaning of "red car" is composed just of the meaning of "red" and the meaning of "car". To understand the combination "red car", you merely have to know what "red" and what "car" means, and you have to know that "red car" is an intersective phrase.

To put it slightly different, "red car" is compositional because "red" is an expression which "makes a uniform semantic contribution to all the compound expressions in which it is embedded".⁴ It does not matter whether "red" is combined with "car" or with other nouns, such as "table" or "book", the meaning which "red" contributes to the meaning of these com-

pounds is always the same. All of them contain one and the same meaning of "red", combined in a certain way with the meaning of the noun. Therefore, you can understand what "red car" means even if you never heard that expression before. It suffices that you know, e.g., what "red table" means and that you know what "car" means because "red" in "red car" makes the same semantic contribution as it makes in "red table".

The meanings of these adjective-noun combinations differ only with respect to the part contributed by the noun. The nouns are the only expressions which give rise to a semantic difference. So, you just have to replace the content which the noun "table" contributes to the meaning of "red table" by the content of "car" in order to work out the meaning of "red car". That is all you have to do because the rest remains unchanged: the meaning of "red car" is composed in the same way of the meaning of its constituents as the meaning of "red table". If you know the meaning of "red table", then there is only one thing you have to learn in order to work out the meaning of "red car": you just have to learn what the noun "car" means.

2. *The Standard Argument against Compositionality*

There are several people who think that what I told you so far is false. A more precise examination of these compounds reveals, so they claim, that it is an illusion to think that "red" makes a uniform semantic contribution to the compounds in which it is embedded. Roughly, their argument against compositionality reads as follows.

The conditions under which an adjective-noun compound containing "red" is applicable differ considerably, depending on the noun which is combined with "red". The conditions under which a car can be called red, for example, are quite different from the conditions under which a newspaper or a watermelon can be called red. More precisely, there are different parts of these objects whose redness is decisive for the application of the expression. Sometimes, it is the surface; in other cases, it is the inside. Therefore, the adjective "red" does not make the same semantic contribution to the compounds containing it.⁵

Let us have a closer look at this standard argument. To have a concrete starting point, I shall describe the version which Ran Lahav presented in an article titled 'Against Compositionality'. I take Lahav's variant of the argument because it is the most extended one I know of.⁶

The first step in the argument is just a repetition of some well-known and very general semantic ideas. A common noun, such as "car", has a cer-

tain extension: there is a class of objects which are correctly described by the word "car". The extension of "car" is determined by the *meaning* of that expression. The meaning of "car" determines its extension by providing the conditions under which something can be called a car. If an object satisfies these conditions, then it belongs to the extension of "car". The meaning of "car" tells us, so to speak, what it is for an object to count as a car by giving the *applicability-conditions* for that noun. The meaning of a common noun determines its extension by providing the conditions under which an object is describable by it.

In the second step of the argument, these ideas are applied to composite expressions. If we combine the noun "car" with the adjective "red", we get a compound with additional applicability-conditions. A red car is, of course, a car. An object satisfies the conditions for being called a red car only if it satisfies the conditions for being called a car. So, the applicability-conditions for "car" are contained in the applicability-conditions for "red car". But "red car" involves additional conditions. There are further constraints an object has to meet in order to be describable by "red car". The applicability-conditions for "red car" go beyond the applicability-conditions for "car" because the adjective "red" provides additional constraints.

The same happens when we combine "red" with another noun, such as "watermelon". The applicability-conditions for "watermelon" are contained in the applicability-conditions for "red watermelon". Adding the adjective "red", however, results in additional constraints. Apart from the conditions provided by "watermelon", an object has to satisfy further conditions in order to be called a red watermelon.

In the third step of the argument, the principle of compositionality appears on the scene. It claims that the part which "red" contributes to the meaning of "red watermelon" is identical with the part it contributes to the meaning of "red car". Since the meaning of an expression is, as we have seen, strongly connected with its applicability-conditions, we can transfer that insight to the applicability-conditions of these compounds. Instead of talking of a uniform *semantic* contribution, we are allowed to talk of a uniform contribution to the conditions under which the compounds can be *applied*. We are allowed to do that because an expression makes the same semantic contribution only if it makes the same contribution to the applicability-conditions.

Modified in that way, the principle of compositionality says that the additional constraints which the adjective "red" contributes to the applica-

bility-conditions of the corresponding compounds are always the same. They do not differ from noun to noun. It does not matter whether you combine "red" with "car" or with "watermelon", the applicability-conditions of these compounds differ merely with respect to the part contributed by the noun. The extra-condition which a car has to satisfy in order to be a red car is the same as the extra-condition a watermelon has to satisfy in order to be a red watermelon. A red watermelon differs from a red car only insofar as it is a watermelon instead of being a car. It does not differ with respect to its being a red object. The additional constraints which a watermelon has to meet for being describable by "red watermelon" are identical with the additional constraints a car has to meet for being describable by "red car". Figuratively speaking, if we remove the applicability-conditions provided by the corresponding nouns, what is left is the same. What it is for a watermelon to count as red is the same as what it is for a car to count as red.

In the fourth step of the argument, this consequence is falsified in order to show that one of the initial assumptions, namely the principle of compositionality, must be false. At this point, Lahav (1989, p. 264) offers an impressive variety of observations concerning compounds with the adjective "red". For the sake of simplicity, let us consider only "red car" and "red watermelon". The conditions under which a car can be called red are different from the conditions under which a watermelon can be called red. The constituents whose redness is decisive for the redness of the whole object are not the same. To be describable by "red car", an object must be a car with a red bodywork; to be describable by "red watermelon", it must be a watermelon with a red pulp. On a more general level, for a car to be red, it must be red *outside*; whereas for a watermelon to be red, it must be red *inside*.

These observations, however, seem to be rather devastating for the principle of compositionality. The part which "red" contributes to the applicability-conditions of "red car" and "red watermelon" is apparently not the same because the additional constraint given by it differs from noun to noun. The consequence is that the adjective "red" makes different semantic contributions to the compounds in which it is embedded because it provides different applicability-conditions. The meaning of such compounds as "red car" and "red watermelon" is not compositional because there is an extra-element in it which differs from one noun to the other.⁷

Knowing the meaning of the combination "red car" and the noun "watermelon" is, therefore, not sufficient for knowing the meaning of "red water-

melon". If you know under which conditions an object can be called a red car, and if you know under which conditions something can be called a watermelon, then you do not automatically know under which conditions an object can be described as a red watermelon. On the contrary, if you apply what you know about the meaning of "red car" to the compound "red watermelon", then you will think that "red watermelon" is applicable only to watermelons which are red outside. And that is false. Semantic competence with respect to "red car" and "watermelon" can easily lead to incompetence with respect to "red watermelon". Competence in the one case is, at least, compatible with incompetence in the other case.

3. Fodor's Account

Is there any way out for a friend of compositionality? There are different ways in which one might try to protect compositionality against the standard argument. Here, however, I want to discuss only Jerry Fodor's reply. It is not too far-fetched to call Fodor a real compositionality addict. Hence, you can imagine that he has a say in that topic.⁸

Fodor's reply is rather radical. From the beginning, it does not even accept that the inside/outside feature is *semantically* relevant. In Fodor's view, it is just as little a part of the *meaning* of "red car" that this expression applies to cars which are red outside as it is part of the *meaning* of "red watermelon" that it applies to watermelons which are red inside. Therefore, there is no threat for compositionality. We need not deliberate where the inside/outside feature comes from because it is not contained in the meaning of the compounds.

According to Fodor (1998b), the adjective "red" is on a par with so-called *relative* adjectives, such as "large". A large ant is surely not as large as a large elephant. Large elephants are much larger than large ants. So, what counts as large in the case of an ant is different from what counts as large in the case of an elephant. The conditions which an *ant* has to satisfy in order to be called large are not the same as the constraints an *elephant* has to meet in order to be called large. Briefly, the applicability-conditions provided by the adjective "large" are different from noun to noun. Hence, in this respect, there is a striking correspondence between "large ant" and "large elephant" on the one hand and "red car" and "red watermelon" on the other hand.

Does that prove, however, that these compounds are non-compositional? No, Fodor says, because these specific applicability-conditions do not

belong to their *semantics*. The meanings of these expressions provide less specific conditions for their application. Of course, whether an ant is large depends on the size of ants, whereas whether an elephant is large depends on the size of elephants. A large elephant is large in comparison with other elephants, whereas a large ant is large in comparison with other ants. In a nutshell, a large elephant is large for elephants, and a large ant is large for ants. But that fact does not threaten compositionality because the meaning of these compounds does not go beyond it.

In Fodor's view, "large ant" means nothing else than *ant which is large for ants*, and "large elephant" just means *elephant which is large for elephants*.⁹ That's it. The meaning of "large elephant" is limited to the information that a large elephant is large for elephants. To know the meaning of "large elephant" is merely to know that this expression applies to elephants which are large in comparison with other elephants. The meaning of "large elephant" does *not* tell us what *counts* as being large in the case of an elephant. It does not tell us which size an elephant must have in order to be describable by "large elephant". This is an empirical, or, more exactly, an encyclopedic, matter which is dependent on how large full-grown elephants actually are. If elephants were much smaller than they are in fact, then the compound "large elephant" would have another extension although the meaning of "large elephant" would remain the same. To vary a passage of Fodor's, what counts as a large elephant depends, not just on matters of meaning, but also on what size elephants actually come in. How could you expect semantics to know which size elephants actually come in? Do you think that semantics runs a zoo or a game reserve?

The same, so Fodor claims, holds in the case of compounds containing the adjective "red". Their meaning is determined in the same way as the meaning of "large elephant". Combining the adjective "red" with the noun "car" leads to an expression which means *car which is red for cars*. Accordingly, "red watermelon" just means *watermelon which is red for watermelons*. That's it. What goes beyond that interpretation is semantically irrelevant because it is not contained in the semantic content of these expressions.

If the meaning of "red watermelon" is in fact limited to *watermelon which is red for watermelons*, then what counts as red in the case of a watermelon is not part of the meaning of that expression. The meaning of "red watermelon" tells us merely that a red watermelon is red for watermelons. It does not contain information about what counts as being red in the case of watermelons. Knowing the meaning of "red watermelon" consists in nothing else than knowing that this expression applies to watermelons

which are red for a watermelon. In particular, its meaning does not tell us which parts of a watermelon are relevant for its colour. It does not tell us whether something is red for a watermelon if it is red inside or outside.

Interpreted in Fodor's way, such compounds are in conformity with the principle of compositionality. There is a uniform function which determines their meaning on the basis of the meaning of their parts. It is irrelevant whether you feed into that function the meaning of "large" and "elephant" or the meaning of "red" and "watermelon", the function always provides a complex meaning with the same structure. You always get a semantic content of the form *[noun] which is [adjective] for a [noun]*.

The phrase "which is such-and-such for a so-and-so" has the function of making explicit the mode in which the constituents are combined. Fodor's interpretation of these compounds indicates, so to speak, their logical deep structure. It shows how the adjective is logically combined with the noun. This mode of combination, together with the meaning of the constituents, determines the meaning of the compound. Hence, a person who knows, e.g., the meaning of "large" and "elephant" and their mode of combination in "large elephant" should have no difficulty in figuring out the meaning of "large elephant".

4. *The Pros and Cons of Fodor's Account*

Fodor's account is very appealing because it is simple and elegant. Moreover, there is a further argument which supports it. It reads as follows.

It is true that large elephants have a size of at least one meter. An elephant which is smaller than one meter is not a large elephant. But do you believe that this is an *analytical* truth? Does the sentence "Large elephants have a size of at least one meter" express a truth solely in virtue of its meaning? I think you will agree that it does not because it is not *necessarily* true. A person who asserts that there are large elephants which are smaller than one meter does not contradict herself because it is easily conceivable that full-grown elephants are at most one meter large. In such a possible world, however, it is not true that large elephants have a size of at least one meter. If evolution had created smaller elephants, the elephants which can be called large would not exceed a size of one meter. Therefore, the sentence "Large elephants have a size of at least one meter" does not express a necessary, and hence no analytical, truth. It is not true solely in virtue of its meaning because the meaning of "large elephant" does not lay down that large elephants have to be larger than one meter. It is true as well because of certain biogenetic facts.

Fodor's account catches that intuition. If "large elephant" just means *elephant which is large for elephants*, then its meaning is silent on the actual size of elephants. Its meaning makes it clear that a large elephant has to be large in comparison with other elephants, but it does not provide the relevant standard of comparison. Consequently, the sentence "Large elephants have a size of at least one meter" does not express an analytical truth. The proposition expressed by that sentence does not go beyond the following statement: elephants which are large for elephants have a size of at least one meter. And that is no analytical truth because it is not necessarily true. In other possible worlds, elephants which are large for elephants might *not* exceed a size of one meter. A person who asserts that there are large elephants which are smaller than one meter does not contradict herself because the meaning of "large elephant" does not lay down that large elephants have to be larger than one meter.

Moreover, we come across a similar insight when we return to the compound "red watermelon". By saying things like "Red watermelons are red inside", Lahav wants to tell us something about the *meaning* of that expression. This observation is not meant as a *biological* remark about red watermelons, but as a *semantical* remark. The fact that red watermelons are red inside is not to be taken as a contingent biological fact. It is to be taken as an analytical fact. The meaning of "red watermelon", so we have to understand Lahav, determines that red watermelons are red inside. A red watermelon which is not red inside is an analytical impossibility.

But if we understand Lahav's remarks in that way, most of them come out false. Is it really true that the expression "red watermelon" is correctly applied only to watermelons which are red inside? Does a speaker say something false when he calls a watermelon red whose pulp is, say, yellow? No, he does not! If the watermelon has a red skin, then what he says might be true.

To be sure, there are no watermelons with a red skin up to now (as far as I know). But that is a contingent fact. It is possible that some passionate melon farmers will cultivate watermelons with a red skin and a yellow pulp. Such a watermelon will not satisfy Lahav's conditions for the application of "red watermelon" because it is not red inside. Nevertheless, it can be called a red watermelon. The meaning of the expression "red watermelon" does not forbid us to correctly describe such a fruit by "red watermelon". Let us assume we are involved in a discussion about beautifully coloured fruits. I point to one of these brand-new melons, saying thereby, 'I like that one. It's a nice red watermelon.' Do I express a falsehood just

because the melon in question is not red inside? No, what I say by "It's a red watermelon" is true. There is no semantical rule to the effect that the extension of "red watermelon" contains only melons which are red inside. The truth that red watermelons are red inside is a contingent truth. It is not an analytical truth. A person does not contradict herself by saying "There are red watermelons with a yellow pulp".

All in all, there is a lot to be said for Fodor's account. Apparently, his interpretation of the compounds in question helps us to get rid of the superfluous ballast which Lahav smuggles into their content. "Large elephant" just means *elephant which is large for elephants*, and "red watermelon" just means *watermelon which is red for watermelons*. If we know that, then we know the meaning of these expressions. What counts as being large in the case of an elephant and what counts as being red in the case of a watermelon is a question which goes beyond their semantic content. In a word, what counts in each case is *semantically underdetermined* (cf. Travis 1981, Part 1; 1996, pp. 454-456).

But there is also a strong case against Fodor's proposal. According to that case, Fodor's interpretation of these compounds leads to counterintuitive consequences because it is not rich enough. A person who is aware of nothing else than what Fodor puts into the semantic content of these expressions will not pass a standard test for understanding.

We have different criteria at hand for determining whether a person knows the meaning of an expression or not (cf. Kühne 1983, Ch. 5, § 3). One of them is that she can provide a correct explanation of the expression. If we ask Susan what "bachelor" means, and she answers "A bachelor is an unmarried man", then we have a good reason to assume that she knows the meaning of "bachelor".

But it is anything but obvious that a person who knows nothing else than what Fodor requires of a competent speaker will pass that test. Let us assume we ask Susan what "large elephant" means. She answers along Fodor's lines, 'A large elephant is an elephant which is large for elephants'. This might be a better answer than the trivial "A large elephant is an elephant which is large". But I am not sure whether we would accept it as an adequate explanation of what "large elephant" means.

However, that is not the only criterion we use. There is a further one which also takes up a central position in ascribing understanding of an expression. It says, roughly, that a competent speaker should be able to apply the expression successfully. There are situations in which she should be ready to apply it; and there are situations in which she should deny that it

is applicable. We can test the semantic competence of a person by confronting her with certain objects and asking for whether the expression in question is appropriate.

Such a test is convincing, of course, only under certain conditions. In the first place, we have to ensure that the person knows enough about the object in order to come to a well-grounded decision. If we point to an elephant far away on the horizon and ask Susan whether it is a large elephant, then we should not be surprised when she does not know what to say. She does not show thereby that she misunderstands the expression "large elephant". Her reaction does not reveal *semantic* incompetence; it is due to missing information about the object in question.

Secondly, we should select only paradigms. The adjective "large" is vague, and it allows for gradation. An elephant can be more or less large, and it can belong to the grey area in which it is neither clear that it is large nor that it is not large. A fair test should start with clear cases. That is, if we want to find out whether Susan knows the meaning of "large elephant", we should confront her, first of all, with elephants which are clearly large and elephants which are clearly not large. Less paradigmatic cases should be included not until she passed that basic test.

The crucial question is whether a person who knows nothing else than what Fodor puts into the semantic content of "large elephant" will pass that test. Is the knowledge which Fodor takes to be sufficient for semantic competence sufficient for a successful performance? The answer is that it is not because such a person may have an inadequate standard of comparison.

Consider Susan who knows that a large elephant is large for elephants. She knows that, in order to determine whether a particular elephant is large, it has to be compared in size with other elephants. Unfortunately, she has the false standard of comparison because, up to now, she has encountered only small elephants. Now, we confront her with Jumbo and ask her whether it is a large elephant. Actually, Jumbo is a paradigm of an elephant which is *not* large. Moreover, Susan has no difficulties in perceiving the size of Jumbo. Nonetheless, Susan answers our question in the affirmative. She thinks that the compound "large elephant" can be applied to Jumbo because, according to *her* standard of comparison, it *is* large. Although Jumbo is a paradigm of an elephant which is not large, Susan is ready to apply the expression "large elephant" because Jumbo is large in comparison with the elephants she knows of.

According to the criterion just-mentioned, this is a strong indication to the fact that Susan does not know the meaning of "large elephant". She is

not able to apply that expression successfully because she does not recognize what counts as a large elephant. She knows that a large elephant must be large for elephants. But this knowledge is not sufficient for passing our test. To be sufficient, it must be backed up by an adequate standard of comparison. Therefore, Fodor's account of the meaning of "large elephant" seems to face a serious difficulty. It is in conflict with a very common and central criterion for understanding an expression. In a word, what Fodor regards as semantic competence is not sufficient for semantic competence according to one of our standard tests.

Things get even worse when we consider a further complication which Fodor missed. I agree completely with Fodor's general idea: the semantics of "red watermelon" has much in common with the semantics of "large elephant". Nevertheless, Fodor has overlooked an important correspondence which makes his *particular* account of their semantics even more problematic.

There are two dimensions which a person has to take into account when she wants to apply a compound containing the adjective "red". You can see that very clearly in the case of "red hair". First of all, as we know already, you have to think about the question which *constituents* of the object in question must be red. But that is not the only dimension which is relevant. Red objects differ also with respect to the *shades of red* which the constituents can have. Red hair, for example, is typically not as red as a beetroot, but reddish brown. In order to apply the compound "red hair" successfully, I do not only have to know that it is the surface which counts; I also have to know in which shades of red hair comes in. If I falsely think that red hair is as red as a beetroot, then I will deny that the compound "red hair" is applicable even when I am confronted with a paradigm of red hair.

Fodor, however, has overlooked that "large" behaves in the same way. There are also two dimensions. Let us assume Susan just knows that "large elephant" applies to elephants which are large for elephants. Fodor seems to think that there is only one additional information Susan has to acquire in order to apply "large elephant" successfully: she has to learn how large elephants can be. This corresponds to the second dimension in the case of "red hair". Knowing how large elephants can be corresponds to knowing how red hair can be. But there is also something which corresponds to the first dimension involved in "red hair". There is a counterpart to knowing which constituents are relevant.

In the case of "large", you have to know whether it is *height* or *length* or *width*, or perhaps a combination of them, which is at issue. You have to

recognize, so to speak, where to set your ruler. And that is an additional dimension about which a speaker à la Fodor has not the faintest idea. If Susan just knows that a large elephant is large for elephants, then she does not even know where to set the ruler. Fodor's explanation of the expression "large elephant" leaves it open to her whether it is a certain height, length, width or a combination of them which counts as large in the case of elephants. The bare knowledge that a large elephant is large for elephants excludes neither of these possibilities. It allows for too many standards of comparison.

So, it is even more clear that Susan will not pass our test for understanding. The knowledge she has is not sufficient for applying the compound "large elephant" successfully. For obvious reasons, the same holds for "red watermelon". Therefore, Fodor's account does not seem to be on the right track. Knowledge of what he puts into the semantic content of "large elephant" and "red watermelon" does not suffice for a satisfactory performance.

5. Conclusion

What is the moral we have to draw from these considerations? To return to the title of my contribution, do compounds such as "red watermelon" and "large elephant" provide a case against compositionality? Or is their meaning, for all that, determined by the meanings of their constituents and their mode of combination? In the introduction, I have announced already that I cannot help but leave open the answer to these questions.

On the one hand, there are good arguments for Fodor's account. In the first place, knowledge about how large full-grown elephants can be does not seem to be semantic knowledge. It is encyclopedic knowledge, hence that information should not be included in the semantic content of "large elephant". Secondly, the sentence "Large elephants have a size of at least one meter" seems to be just as little an analytical truth as "Red watermelons are red inside". Therefore, again, such information should not be included in the meaning of these expressions. Thirdly, the meaning of "red watermelon" does not lay down whether it is the inside or the surface which counts. What counts is semantically underdetermined; it depends on the context of utterance. Hence, again, that information should not be included in the semantic content of "red watermelon".

All in all, it is arguable that the meaning of such compounds is relatively undemanding and that it does not go beyond Fodor's interpretation.

If this is correct, then there is no threat for compositionality. Their meaning is, then, composed in a uniform way of the meaning of their parts.

On the other hand, there is an argument against Fodor's account. According to a standard test, a person who knows the meaning of an expression should be able to apply it successfully. In the case of "large elephant", this requires knowing where to set your ruler and knowing how far elephants can extend across the corresponding dimension. Hence, it seems that this information should be included in the semantic content of "large elephant".

But if that is true, then there is little hope for compositionality. Knowing, for example, the meaning of "large elephant" and the meaning of "ant", is, then, not sufficient for understanding "large ant". For in order to understand "large ant", you have to know how large ants can be. But neither the meaning of "large elephant" nor the meaning of "ant" tells you which size an ant actually can have.

At the moment, I do not know how to get out of that dilemma.¹⁰ Probably, it is based on different conceptions of semantics. We have to think about which explanatory tasks a semantic theory has to see to. I am not sure what Fodor demands of semantics, but it is obvious that his conception is relatively modest. It does not provide the basis on which a person applies an expression successfully. A semantics à la Fodor must be backed up by further considerations in order to explain why some people show a satisfactory performance whereas other people do not. In a word, Fodor's conception is only loosely connected with the way in which competent speakers *use* an expression.

But there is another conception of semantics. According to it, meaning is strongly connected with the way in which competent speakers use an expression because knowing the meaning of an expression is strongly connected with the ability to apply it in a satisfactory way. That is, I think, the everyday notion because it is the notion behind some of our usual tests for understanding. Such a conception cannot be content with Fodor's proposal because that proposal requires too little of a person who understands the corresponding expressions.

Moreover, it might be helpful to take into account a fact which only few people bear in mind. It is the simple fact that understanding allows for gradation (cf. Dascal 1981, p. 335f.). You can know the meaning of an expression more or less; and perhaps you can also improve your understanding by acquiring knowledge which is usually thought of as encyclopedic. So, it might be that Fodor's interpretation merely describes what it is to understand those compounds *partially* or *rudimentarily*. He offers, so to speak,

the *core* of their semantic content which is compositionally determined. The standard test for understanding, on the other hand, is a test for *complete* or, at least, *more perfect* understanding. It goes beyond the semantic core by appealing to information which is not compositionally determined.¹¹ But that is an idea which requires further investigation.¹²

Notes

- ¹ Cf., e.g., Bennett 1976, p. 16f.; Fodor & Pylyshyn 1988, p. 12; Goschke & Koppelberg 1991, p. 138; Lyons 1995, p. 46; Kamp & Partee 1995, p. 135; Platts 1979, p. 70; and Schiffer 1987, p. 179.
- ² Lyons 1995, p. 112; cf. p. 204. Manfred Pinkal (1985, p. 33) claims, along the same lines, that one cannot pursue a serious semantics without taking in account such a principle.
- ³ Cf. Bennett 1976, p. 155; Fodor 1994, p. 106f.; Fodor 1998a, pp. 94-96; Fodor & Lepore 1996, pp. 254-258; Johnson-Laird 1983, p. 169; Lakoff 1987, p. 146; Murphy 1988, pp. 529, 558; Lyons 1995, pp. 206, 228; Platts 1979, p. 43; Thagard 1992, p. 51; and Ziff 1960, p. 61.
- ⁴ Lahav 1989, p. 261. For similar formulations of the principle of compositionality, cf. Fodor & Pylyshyn 1988, p. 42; and Goschke & Koppelberg 1991, p. 138.
- ⁵ Searle (1979; 1980; 1983, Ch. 5) has pointed to analogous difficulties with respect to verbs, such as "open" and "cut", and whole sentences, such as the philosophical paradigm "The cat is on the mat".
- ⁶ For a rather short version, cf. Goschke & Koppelberg 1991, p. 144.
- ⁷ As some cognitive psychologists say, there are *emergent* features in the meaning of these adjective-noun combinations (cf. Rips 1995, p. 92f.).
- ⁸ Nevertheless, he hides it in short remarks and one or two footnotes.
- ⁹ Actually, Fodor says that "(is a) large elephant" means (*is*) *large for elephants* (not (*is an*) *elephant which is large for elephants*). But there is no real difference between these ways of explaining the meaning of such compounds because, as Platts (1979, p. 184) has noticed, being *F* for a *G* already implies being a *G*. However, it is easy to overlook that implication; hence I prefer to use formulations of the type "*G* which is *F* for a *G*" because they make it explicit.
- ¹⁰ Have a look, however, at Díez's comment where he solves the dilemma in a conclusive way.
- ¹¹ This proposal approaches a linguistic theory called *two-level semantics* (cf. Bierwisch 1983, Bierwisch & Lang 1989) and the distinction between *defining* and *characteristic* features (the *core* of a concept and connected *identification procedures*) offered by some cognitive psychologists (cf. Hampton 1988, p. 12f.; Osherson & Smith 1982, p. 317; Rips 1995, p. 79ff.).
- ¹² Many thanks to Frank Kannetzky, Christian Plunze, Richard Raatzsch and Mark Textor for discussions on an earlier draft and to José A. Díez and the participants of the

workshop for their very helpful comments on my talk. As you will see, following up Díez's suggestions would lead to a different paper, not the one to which he replied at the workshop.

BIBLIOGRAPHY

- Bennett, J.: 1976, *Linguistic Behaviour*, Cambridge.
- Bierwisch, M.: 1983, 'Semantische und konzeptuelle Repräsentation lexikalischer Einheiten', in W. Motsch & J. Kunze (ed.): *Untersuchungen zur Semantik (studia grammatica 22)*, Berlin, pp. 61-99.
- Bierwisch, M. & Lang, E. (eds.): 1989, *Dimensional Adjectives. Grammatical Structure and Conceptual Interpretation*, Berlin.
- Dascal, M.: 1981, 'Strategies of Understanding', in R. Posner (ed.): *Meaning and Understanding*, Berlin & New York, pp. 327-352.
- Fodor, J.: 1994, 'Concepts: A Potboiler', *Cognition* 50, 95-113.
- Fodor, J.: 1998a, *Concepts. Where Cognitive Science Went Wrong*, Oxford.
- Fodor, J.: 1998b, 'There Are No Recognitional Concepts; Not Even RED', online-paper, <http://www.nyu.edu/gsas/dept/philo/courses/concepts/FodorCh4.htm>.
- Fodor, J. & Lepore, E.: 1996, 'The Red Herring and the Pet Fish: Why Concepts still can't be Prototypes', *Cognition* 58, 253-270.
- Fodor, J. & Pylyshyn, Z.W.: 1988, 'Connectionism and Cognitive Architecture: A Critical Analysis', *Cognition* 28, 3-71.
- Goschke, T. & Koppelberg, D.: 1991, 'The Concept of Representation and the Representation of Concepts in Connectionist Models', in W. Ramsey, S.P. Stich & D.E. Rumelhart (ed.): *Philosophy and Connectionist Theory*, Hillsdale/NJ, pp. 129-161.
- Hampton, J.A.: 1988, 'Overextension of Conjunctive Concepts: Evidence for a Unitary Model of Concept Typicality and Class Inclusion', *Journal of Experimental Psychology: Learning, Memory, and Cognition* 14, 12-32.
- Johnson-Laird, P.N.: 1983, *Mental Models. Towards a Cognitive Science of Language, Inference, and Consciousness*, Cambridge.
- Kamp, H. & Partee, B.: 1995, 'Prototype Theory and Compositionality', *Cognition* 57, 129-191.
- Künnc, W.: 1983, *Abstrakte Gegenstände. Semantik und Ontologie*, Frankfurt/M.
- Lahav, R.: 1989, 'Against Compositionality. The Case of Adjectives', *Philosophical Studies* 57, 261-279.
- Lakoff, G.: 1987, *Women, Fire, and Dangerous Things*, Chicago.
- Lyons, J.: 1995, *Linguistic Semantics. An Introduction*, Cambridge.
- Murphy, G.L.: 1988, 'Comprehending Complex Concepts', *Cognitive Science* 12, 529-562.
- Osherson, D.N. & Smith, E.E.: 1982, 'Gradedness and Conceptual Combination', *Cognition* 12, 299-318.
- Pinkal, M.: 1985, *Logik und Lexikon. Die Semantik des Unbestimmten*, Berlin & New York.
- Platts, M.: 1979, *Ways of Meaning. An Introduction to a Philosophy of Language*, London.
- Rips, L.J.: 1995, 'The Current Status of Research on Concept Combination', *Mind & Language* 10, 72-104.
- Schiffer, S. 1987, *Remnants of Meaning*, Cambridge/Mass.
- Searle, J.R. 1979, 'Literal Meaning', in *Expression and Meaning*, Cambridge, pp. 117-136.

- Searle, J.R.: 1980, 'The Background of Meaning', in J.R. Searle, F. Kiefer & M. Bierwisch (ed.): *Speech Act Theory and Pragmatics*, Dordrecht, pp. 221-232.
- Searle, J.R.: 1983, *Intentionality. An Essay in the Philosophy of Mind*, Cambridge.
- Thagard, P.: 1992, *Conceptual Revolutions*, Princeton/NJ.
- Travis, Ch.: 1981, *The True and the False. The Domain of the Pragmatic*, Amsterdam.
- Travis, Ch.: 1996, 'Meaning's Role in Truth', *Mind* 105, 451-466.
- Ziff, P.: 1960, *Semantic Analysis*, Ithaca/NY.

Mark Siebel has done his PhD in philosophy at the University of Hamburg in 1998. He is an assistant in the research group *Kommunikatives Verstehen* (*Communicative Understanding*) at the University of Leipzig. His areas of interest are philosophy of language, communication theory, philosophy of mind, epistemology and Bernard Bolzano. He has published the books *Der Begriff der Ableitbarkeit bei Bolzano* (*Bolzano's Concept of Derivability*, 1996) and *Erinnerung, Wahrnehmung, Wissen* (*Memory, Perception, Knowledge*, 2000).