

ADJECTIVES IN CONTEXT

Zoltán Gendler Szabó
The Sage School of Philosophy
221 Goldwin Smith Hall
Cornell University
Ithaca, NY 14853-3201
e-mail: zs@cornell.edu

In R. Harnish and I. Kenesei eds.,
Perspectives on Semantics, Pragmatics, and Discourse.
Amsterdam: John Benjamins, 2001, pp. 119-46.

0. Abstract

In this paper, I argue that although the behavior of adjectives in context poses a serious challenge to the principle of compositionality of content, in the end such considerations do not defeat the principle. The first two sections are devoted to the precise statement of the challenge; the rest of the paper presents a semantic analysis of a large class of adjectives that provides a satisfactory answer to it. In section 1, I formulate the context thesis, according to which the content of a complex expression depends on the context of its utterance only insofar as the contents of its constituents do. If the context thesis is false, the content of some complex expression is not compositionally determined. In section 2, using an example due to Charles Travis, I construct an objection to the context thesis based on the behavior of the adjective ‘green’. In section 3 and 4, I look at some of the difficulties surrounding the semantics of ‘good’, which provide the motivation for the thesis that most adjectives are contextually incomplete one-place predicates. In section 5, I discuss how ‘green’ and other color adjectives can be treated within such a semantic theory. Since this theory is compatible with the context thesis, the objection against the compositionality of content loses its force.

1. The Context Thesis

One of the fundamental assumptions of contemporary semantic theory is the principle of compositionality, the thesis that the meaning of a complex expression is determined by the meanings of its constituents and by manner in which these constituents are combined.

When challenged, defenders of the principle tend to appeal to the undeniable fact that competent speakers routinely understand complex expressions they have never heard before. This fact, they contend, can only be adequately explained if we assume that linguistic competence requires implicit grasp of a recursive system of semantic rules

* This paper is a substantially revised version of chapter 4 of my dissertation. I want to thank the late George Boolos, Richard Cartwright, Robert Stalnaker and Judith Thomson for discussion and advice regarding drafts of that chapter. The final version of the paper benefited from discussions with Tamar

which assign meanings to complex expressions on the basis of the meanings of their constituents and the way those constituents are combined.

But this is much too quick. Nobody denies that the meaning of a complex expression *depends* on the meanings of its constituents and on its structure; the bite of compositionality is that it depends on *nothing else*. The appeal of the principle stems from the conviction that those who can understand a new expression must work out its meaning from facts they knew about the expression antecedently. And of course, the only candidate facts are lexical or syntactic. But this piece of reasoning neglects the possibility that the *context* in which someone encounters a new expression can provide additional clues relevant for interpretation. And if it does, there is no reason to assume that the lexical and syntactic facts *by themselves* fix the meaning of the new expression. So at the very least, attention to context undermines the unreflective assumption that compositionality is a trivial matter.¹

Let us make our terminology more precise. The word ‘meaning’ is used in semantics to refer to all sorts of semantic values; in what follows, I want to focus on a particular semantic value often called ‘content’. The *content* of a declarative sentence is

Szabó Gendler, Harold Hodes, Jeff King, and Jason Stanley. I also wish to thank two anonymous referees for this volume whose comments led to further revisions.

¹ As I understand the principle of compositionality, it says that the meaning of complex expressions – phrases, clauses and sentences – depends on two kinds of facts. Syntax tells us what the lexical constituents – words and smaller morphemes – occurring in a complex expression are and how they are combined; and the lexicon tells us what those constituents mean. According to the principle, such syntactic and semantic features of a complex expression are jointly sufficient to determine the meaning of that expression. But what does it take for certain features of an expression to *determine* what that expression means? It is customary in semantics to assume that the relevant determination relation is simply functional dependence. I believe this interpretation of the principle of compositionality is unreasonably weak. In Szabó (2000) I argue that the principle is best understood as the claim that there is *a single function across all possible human languages* that assigns the meanings of complex expressions to the meanings of their constituents and their syntactic manner of composition. This, however, will not matter for the purposes of the present discussion.

the proposition it expresses and the contents of sub-sentential expressions are the semantic contributions of such expressions to propositions expressed by declarative sentences within which they occur as constituents. There is no general agreement concerning the nature of propositions; the only thing I will assume about them is that they determine truth-conditions, that is, that if two sentences express the same proposition they cannot differ in truth-value. Semantic theories typically aim at ascribing content to complex expressions in a compositional fashion, hence they are committed to the principle of compositionality of content.²

I take the *context* of utterance to be 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 in question was made, and perhaps much more. There are several important foundational questions about contexts, but for the purposes of this paper I wish to remain more or less neutral with regard to these.³ I will consider the context of an utterance

² One might object to calling the content of an expression its “meaning”. If we think that the meaning of an expression is a feature knowledge of which guarantees understanding of the expression then the proposition expressed by a sentence is not a good candidate for being the meaning of that sentence. (Intuitively, we all know what ‘I am tired’ means but we don’t know what proposition it expresses on a given occasion unless we know who made the utterance.) Perhaps, we should follow David Kaplan and think of meanings as *characters*, i.e. functions from contexts to contents; cf. Kaplan (1977). Be that as it may, compositionality of content is at least as important in semantics as compositionality of meaning proper. Just as we can speak about understanding *expressions*, we can also speak about understanding *utterances* and the latter is presumably tied to grasping the proposition expressed by the utterance. So, insofar that it is plausible that we can understand not only sentences we never hear before but also particular utterances of such sentences, we have some intuitive support for the principle of compositionality of content.

³ There are two major traditions of thinking about context within semantics. There are those who follow Montague (1968) in representing contexts as *indices*, i.e. as an *n*-tuples of speaker, hearer, time, place, a series of salient objects, etc. And there are those who follow Stalnaker (1970) in thinking of contexts as

simply as the sum total of *all* the non-lexical and non-syntactic facts relevant for determining what a speaker conveys by the utterance. Let us say that an expression is *context-dependent* if and only if it can have different contents when uttered in different contexts. A context-dependent complex expression none of whose constituents are context-dependent would then be a counterexample to the principle of compositionality of content. Is there such a counterexample?

There are many cases when we *seem* to be presented with context-dependent expressions, but in fact we are not. Such cases come in two varieties, which I will call *phonetic* and *pragmatic* illusions of context-dependence. Phonetic illusions of context-dependence occur when we have different expressions articulated in the same way.

Consider the following two sequences of sentences:

- (1a) I should stop digging this hole in the bottom of this ship, but I cannot resist the temptation. Sinking boats can be entertaining.
- (1b) I should stop watching sea battles on video, but I cannot resist the temptation. Sinking boats can be entertaining.

According to their most plausible interpretations, (1a) is about my activity of sinking boats, while (1b) is about boats that are sinking. So it seems that the content of ‘Sinking boats can be entertaining’ varies depending on the larger linguistic context in which it is used. But this is not so. The illusion that we have expressions with multiple contents arises because our ordinary *notation* is ambiguous: ‘Sinking boats can be entertaining’ refers to two different sentences.⁴ The justification for the claim that we are dealing with

information states which can be represented as sets of possible worlds (or sets of situations) compatible with the shared knowledge of the participants in the conversation.

⁴ I use the following conventions of quotation. An expression *e* between single quotes is a (possibly ambiguous) name for *e*. I also use the iterated single quote to name quotation names of expressions. Double quotes differ from iterated single quotes in appearance: in the case of iterated single quotes I leave a space between the inverted commas, in the case of double quotes I do not. Double quotes are strictly

two different expressions here is syntactic: the linear order of ‘sinking’ and ‘boats’ disguises glaring structural differences.

Pragmatic illusions of context-dependence may arise when utterances of an expression *convey* different things in different contexts. From this, one might be tempted to conclude immediately that the expression *says* different things in different contexts.

But this does not follow. Consider these sequences:

- (2a) I could not solve the crossword puzzle, even after I wasted the whole weekend on it. The puzzle was extremely difficult.
- (2b) I solved the crossword puzzle in five minutes while cooking dinner and listening to the radio. The puzzle was extremely difficult.

It is quite plausible to take ‘The puzzle was extremely difficult’ as ironical in (2b), but not in (2a). Still, there is no good reason to postulate that this sentence expresses different propositions in these two cases; its content is invariably the proposition that the puzzle was extremely difficult. But since this proposition is in conflict with the content of the first sentence of (2b), a sensible hearer will conclude that in uttering the second sentence of (2b) the speaker did not mean to convey what he literally said. We have strong reasons to think that irony is a pragmatic matter: there is no foolproof conventional mark indicating that an utterance is ironical. Moreover, we couldn’t even introduce such a conventional sign. For suppose we agreed that, say, a winking of the left eye signals that the utterance must be interpreted ironically; those who wanted to be *genuinely* ironical would certainly refrain from winking in making their remarks.

speaking not quotation devices: the expression between the double quotes is used, not mentioned. (Or perhaps, partly used, partly mentioned.) The following examples are illustrative:
(i) ‘Cambridge’ refers (ambiguously) to the cities Cambridge, England and Cambridge, Massachusetts.
(ii) ‘ ‘Cambridge’ ’ refers to an ambiguous quotation name which is used in (i).
(iii) According to a friend of mine, “Cambridge is a reasonably lively place.” He probably has Cambridge, Massachusetts in mind.

Despite the prevalence of phonetic and pragmatic illusions, cases of genuine context-dependence are not hard to come by. Consider (3) and (4), for example:

(3) Some basketball players are giants.

(4) I am now at the North Pole.

What (3) and (4) say is fixed in part by the context of utterance. If a giant is supposed to be a “legendary being of great stature and strength and of more than mortal but less than godlike power” then (3) is false, but if it is merely “a living being of great size”, it is true.⁵ Since the fact that ‘giant’ can have different contents is clearly marked in the lexicon, (3) is not a likely candidate for a pragmatic illusion of context-dependence. Saying that there are two words in English – ‘giant₁’ and ‘giant₂’ – and that consequently we are dealing with a phonetic illusion of context-dependence in (3) would also be implausible.⁶ Neither semantic intuition, nor etymology would support such a claim. A speaker who never heard ‘giant’ be used in referring to people of extraordinary size would probably still be able to figure out how to understand (3); by contrast a speaker who doesn’t know that ‘bank’ can be used to refer to the edges of rivers will not know what to make of ‘I closed my account in the bank.’ ‘Giant’ is a *polysemous* expression, i.e. ‘giant’ refers unambiguously to ‘giant’, but this lexical item has multiple closely associated contents.⁷

These conventions are not supposed to be taken as an analysis of how quotation is ordinarily used in English. I have grave doubts about the coherence of our ordinary use quotation marks.

⁵ The definitions are from the *Webster Dictionary*.

⁶ Here and throughout the paper I use the Courier font with or without indices to indicate disambiguation.

⁷ Unfortunately, there is no terminological unity in linguistics concerning types of ambiguity. I prefer the following definitions. A linguistic type (a phoneme or a sequence of phonemes) is *ambiguous* if it has more than one content. If an ambiguous linguistic type stands for a single expression, it is *polysemous*; otherwise it is *homonymous*.

The case of (4) is even more straightforward: different people can say different things at different times uttering this sentence. I have never been in a position to use it to express a truth, but I might be in the future and others have been in the past. ‘I’ and ‘now’ are indexical expressions; their contents are fixed in accordance with a linguistic rule relative to the context of use. And the context-dependence of indexicals spreads to the sentences that contain them as constituents.

Neither polysemy nor indexicality is a threat to compositionality: (3) and (4) have multiple contents *because* they contain constituents whose contents depend on the context. To find a compelling counterexample to compositionality we would need a complex expression whose content changes with the context in which it is used even though the contents of all of its constituent expressions remain fixed. Whether or not there is such a counterexample, coming up with one is not a trivial matter. I will call the claim that there is no such counterexample the context thesis:

Context Thesis: The content of an expression depends on context only insofar as the contents of its constituents do.

According to the context thesis the context-dependency of a complex expression must be indirect. Once one has fixed what content the simple constituents have in a given context, the way in which the content of the complex expression depends on the context is also fixed. In the next section, I will consider an argument against the context thesis, based on the semantic behavior of certain complex nominal expressions containing adjectives.

2. The Color of a Painted Leaf

Charles Travis has argued that many different things might be said with the same words, even if the subject matter of the utterances remains the same:

As an arbitrary example, consider the words ‘The leaf is green’, speaking of a given leaf, and its condition at a given time, used so as to mean what they do mean in English. How many distinct things might be said in words with all that true of them? Many. That emerges when we note that one might speak either truth or falsity in such words, if the leaf is the right way. Suppose a Japanese maple leaf, turned brown, was painted green for a decoration. In sorting leaves by colour, one might truly call this one green. In describing leaves to help identify their species, it might, for all the paint, be false to call it that. So words may have all the stipulated features while saying something true, but also while saying something false. Nothing about what it is to be green decides whether the colour of a thing is the way it is with, or the way it is without the paint. What being green is is compatible with speaking either truth or falsity in calling the leaf green. For all that, the painted leaf is as it is sometimes, but not other times, said to be in calling it green.⁸

What the example shows is subject to debate. Let me start with what is relatively uncontroversial. First, the truth-value of an utterance of ‘The leaf is green’ made about the painted leaf depends on the context. Roughly, the utterance is true if it is made when one is sorting leaves for decoration, and false if it is made when one is trying to identify the species of the leaf. Second, the sequence of phonemes uttered in the two scenarios stand for the same unique sentence, i.e. ‘The leaf is green’ refers unambiguously to the sentence ‘The leaf is green’. Third, the words ‘leaf’ and ‘green’ are used in both cases “so as to mean what they do mean in English”, which is presumably something fixed. Consequently, we have a single sentence which says different things in different contexts, even though the sentence is neither syntactically nor lexically ambiguous.

⁸ Travis (1994), pp. 171 - 2.

Is this a challenge to the context thesis? It depends. ‘Meaning’ is an elusive phrase and it is not entirely clear how Travis intends to use it here. If by ‘meaning’ we mean the conventions of use for an expression knowledge of which guarantees understanding then the fact that sentences with a fixed meaning can say different things is no surprise. Any sentence containing an indexical has a single meaning in this sense even though it can express a plethora of different propositions. But suppose ‘meaning’ is taken to mean the contribution an expression makes to the proposition expressed by sentences containing the expression. Then the claim that ‘The leaf is green’ can express different propositions even though the words this sentence is constructed from have a fixed meaning is in direct conflict with the context thesis.

The obvious reply to this challenge is to reject the claim that the constituent expressions have identical contents in the two utterances. One idea would be to say that this is a case of referential indeterminacy: ‘the leaf’ can denote the leaf with or without the paint. There is some plausibility to this idea, given that we would be strongly inclined to call both of those entities leaves. No doubt, this line gives rise to some metaphysical puzzles (How come we have two leaves roughly at the same spatio-temporal location? Do we really create a new leaf in painting one? If we remove the paint is the leaf with paint destroyed?), but the solution is implausible independently of these puzzles. For in some contexts ‘the leaf’ denotes what is below the paint, then presumably in those contexts ‘The leaf is unpainted’ is true. But this is unacceptable: there is *no* context in which one can hold up a painted leaf and say truthfully ‘The leaf is unpainted’.

Clearly, it is better to locate the context-dependence in the adjective. One might say that there is no such thing as *the* content of ‘green’, just as there is no such thing as

the content of ‘giant’ or *the* content of ‘I’. What it is to be green varies with the context. But there are two problems with this reply. The first is that according to Travis the example was merely illustrative, and the phenomenon has nothing to do with specificities of ‘green’. The second is that ‘green’ does not seem to be a polysemous or indexical word. Let us look at these problems in more detail.

According to Travis, the example of painted leaf can be easily generalized. In fact, he thinks the moral applies to “whatever else words speak of”:

The words ‘is green’, while speaking of being green, may make any of indefinitely many distinct contributions to what is said in words of which they are part. The above variation is illustrative. The same holds of any English predicate. The fact that ‘is green’ speaks of being green does not alone decide what is required for a thing to be as it, on a speaking, says a thing to be. Similarly for whatever else words speak of.⁹

I think this is an overstatement. If one uses the sentence ‘The number is even’, talking about the number four, the sentence expresses a truth. One cannot construct some special scenario where due to some camouflage similar to the painting of the leaf ‘The number is even’ says something false.

There is a trivial way that the example can be generalized to practically all non-mathematical and non-logical cases. Take the sentence ‘The book is a novel’. There are all sorts of borderline instances of novelhood, like longish short-stories, modern epics written in free verse, etc. Let the phrase ‘the book’ be used to denote some such entity. Then there are scenarios where the sentence is true according to the contextual norm, and others where it is false. This is not especially interesting, since it does not threaten the context thesis: intuitively, whenever the content of ‘... is a novel’ is vague in *this* way, it is because the content of ‘novel’ is vague. What is challenging about the example of the

⁹ Travis (1994), p. 172.

painted leaf is that it is independent of problems of vagueness. Even if a brown maple leaf is painted the most paradigmatic green one can imagine, the problem whether the leaf is green remains.

If we neglect vagueness, however, Travis's example is not widely generalizable. There are indeed different scenarios where utterances of the sentence 'The ring is gold' are about the same object, but have different truth-values. If the ring contains silver besides pure gold the utterance can be true in ordinary circumstances and false in a laboratory. But speaking of a ring made of pure gold, 'The ring is gold' expresses a truth, no matter what the speaker thinks, or what the purpose of her statement is.

So, we can conclude that the kind of context-dependency that we see in 'The leaf is green' is *not* universal: 'The number is even' or 'The ring is gold' are rather similar sentences that do not have it. Of course, it is equally obvious that the example is generalizable to many other adjectives, besides 'green'. It is quite easy to come up with parallel cases for 'The apple is red', 'The move is smart', 'The problem is interesting', or 'The soup is good'. The question is then, what is responsible for this limited phenomenon?

This brings us to the second problem concerning the suggestion that 'green' is a context-dependent adjective. One can articulate the intuition behind Travis' example thus: "Of course, the challenge can be countered by saying that 'green' expresses one content in one scenario and another content in the other. But intuitively, this is not the case. What it contributes to a proposition expressed by 'The leaf is green' in either scenario is the same: the property of being green. 'Green' is not a polysemous expression and it is not an indexical either. To stipulate that it is nevertheless context-dependent

requires that we introduce context-dependence of a new, hitherto unknown sort. And this seems like an *ad hoc* maneuver.”

If we look up ‘green’ the *Webster Dictionary*, what we find – alongside the non-literal meaning specifications, which are irrelevant here – is the barely exciting information: something is *green* if it is “of the color green”. This fact is indeed in conflict with both the assumptions that ‘green’ is polysemous and that it is an indexical. If ‘green’ were a polysemous word like ‘giant’, we would expect the dictionary to contain different numbered sub-entries specifying the different contents the word can have on different occasions of its use. If ‘green’ were an indexical like ‘I’, we would expect the dictionary to contain an informative clause which tells us how to select its content in a given occasion of its use. If ‘green’ is context-dependent, its context-dependence is of a different kind.

The question is whether a convincing story can be told about *how* the content of ‘green’ and other adjectives is supposed to depend on the context. If there is no such story, the sentence ‘The leaf is green’ refutes the context thesis and provides a genuine counterexample to the principle of compositionality of content. But I don’t think the situation is as bad as all that. In the rest of this paper, I will present the outlines of an account of the semantics of a large class of context-dependent adjectives. And I will argue that this class includes ‘green’.

3. Problems with ‘Good’

It is a curious fact that contemporary views on the semantics of adjectives evolved from a debate that started in moral philosophy. In their effort to articulate theories about what goodness consists in, philosophers turned to questions about the semantics of ‘good’. ‘Good’ is an adjective with some peculiar characteristics, so any analysis of its content had to say something about the interpretation of adjectives in general as well as about the semantic features that distinguish ‘good’ from less problematic adjectives, like ‘round’ or ‘tall’.

Like most adjectives, ‘good’ occurs in *predicative* as well as *attributive* positions. The former cases suggest that ‘good’ is a predicate, while the latter that it is a predicate-modifier. Which of these indications is to be taken seriously is a question that might affect the way in which we make sense of judgments to the effect that an act, a person, or a state of affairs is good.

There is a certain difficulty about pronouncements that ‘good’ is a predicate, or that it is a predicate-modifier. Categories like ‘noun’, ‘verb’, ‘determiner’, ‘pronoun’, ‘adjective’, or ‘sentence’ are syntactic, and categories like ‘predicate’, ‘connective’, ‘quantifier’, ‘predicate-modifier’, ‘variable’, or ‘formula’ are logical. These logical categories, though they are syntactic categories of certain formalized languages, are certainly not syntactic categories of English.¹⁰ But what is it supposed to mean that an expression of English belongs to a syntactic category of some *other language*?

One natural way to understand such claims is to regard them as being about *logical forms*. I take the logical form of an expression *e* in a natural language *L* to be an

¹⁰ What makes a category syntactic is that we can determine which expressions belong to that category through morpho-syntactic means. Syntactic categories of formal languages can be ‘read off’ the rules which specify the class of well-formed formulas in the language.

expression e' in a suitable formalized language L' which has the same *logical properties* as e . If L' has a clear, well-understood semantics e' can be used to illuminate and explain the inferential behavior of e in L .¹¹ This means that an expression could have more than one logical form. If we can choose among the competing logical forms one that explains the inferential behavior of an expression in the most perspicuous way, we can call this *the* logical form of the expression.

To call 'good' a predicate (or to say that the semantic category for 'good' is that of predicate) is then to say that in the logical form of English sentences containing the word 'good' we always find the same corresponding predicate. Similarly, to call it a predicate-modifier is to claim that in the logical form of English sentences containing 'good' we always find the same corresponding predicate-modifier. These claims can be formulated as (5) and (6):¹²

(5) The logical form of 'good' is the predicate ' $\text{good}(x)$ '

(6) The logical form of 'good' is the predicate-modifier ' $(\text{good}(F))(x)$ '

(Here ' x ' is an individual variable and ' F ' is a predicate variable.)

G. E. Moore claimed in *Principia Ethica* that 'good' is a predicate that expresses a simple property. Certainly, there is a *prima facie* plausibility in the idea that 'good'

¹¹ One might wish to steer between an extremely weak and an extremely strong conception of logical form. According to the weak conception, the only criterion that we might use to evaluate the adequacy of a certain logical form is whether it gets all the inferences right. This would mean that we can assign the same logical form to any two sentences of English that are logically equivalent. But surely, we should not ascribe 'Snow is falling' and 'Snow is falling and snow is self-identical' the same logical form. The strong conception of logical form requires that an expression and its logical forms share all their semantically relevant properties. This would demand that semantic features, like [INANIMATE] be represented in logical form. But it is not clear whether we really want to ascribe different logical forms to 'It is here' and 'She is here'.

¹² I use the Courier font for logical forms, which is the notation introduced in the previous section to indicate disambiguation; cf. fn. 6. Given that we use logical forms to *clarify* and *explain* the logical behavior of natural language expressions, the demand that the languages of logical forms be free of ambiguity seems reasonable. If I had the goal of absolute precision, I would use corner quotes whenever I mention logical forms containing free variables. Since I don't, I won't.

expresses goodness, like ‘hard’ hardness or ‘green’ greenness. If there is no reason to think otherwise, we can suppose that expressions in the same grammatical category have the same *kind* of contents.

A semantic analysis based on Moore’s suggestion regards the predicative occurrences of ‘good’ as paradigm cases and treats sentences containing attributive occurrences as complexes built up from simpler sentences in which there are only predicative occurrences. According to Moore, something is a good book if and only if it is good and it is a book.¹³ One way to implement these ideas is to say that the logical form of (7) is (7a), where the language of logical forms is a formalized extensional first-order language:

(7) Sue is a good dancer (7a) $\text{dancer}(\text{Sue}) \wedge \text{good}(\text{Sue})$

There is an argument due to Peter Geach, which shows that this simple Moorean analysis cannot be the whole story about ‘good’. The argument is based on the observation that inferences of a certain type are incorrectly predicted to be valid by the Moorean analysis. Let us call an inference a *transfer-inference*, just in case it fits the following pattern: its first premise is of the form ‘ a is AN ’, its second premise has an entailment of the form ‘ a in M ’ and its conclusion is of the form ‘ a is AM ’. Moore’s analysis is in trouble with the following simple transfer-inference:¹⁴

(7) Sue is a good dancer	(7a) $\text{dancer}(\text{Sue}) \wedge \text{good}(\text{Sue})$
<u>(8) Sue is a pianist</u>	<u>(8a) $\text{pianist}(\text{Sue})$</u>
(9) Sue is a good pianist	(9a) $\text{pianist}(\text{Sue}) \wedge \text{good}(\text{Sue})$

¹³ One of the many passages where Moore endorses this view is the following: “For ‘good conduct’ is a complex notion: all conduct is not good; for some is certainly bad and some may be indifferent. And on the other hand, other things, beside conduct, may be good; and if they are so, then, ‘good’ denotes some property, that is common to them and conduct ...” Moore (1903), p. 2.

¹⁴ Cf. Geach (1956).

The inference on the right is trivially valid, but according to its most natural interpretation, the inference on the left is invalid; Sue might be a good dancer and a pianist without being a good pianist. The fact that ‘good’ fails to validate certain transfer inferences is regularly referred to in the literature as the *non-transparency* of this adjective.¹⁵ The non-transparency of ‘good’ shows that at least one of the Moorean translations from the English sentences to the logical formalism must be incorrect.

Geach’s suggestion was to give up the Moorean assumption that ‘good’ is a predicate and regard it instead as a predicate-modifier. On this approach, the attributive occurrences are the paradigm cases and the predicative occurrences must be explained away. To give an account of them, one has to assume that they are elliptical attributive occurrences. Thus the sentence ‘These gloves are good’ is usually to be understood as ‘These gloves are good gloves’. Although this is a natural interpretation of the sentence, in certain situations it is incorrect. For example, if the sentence is uttered by someone who wants to clean the windows and is looking for some piece of cloth appropriate for that purpose, the suggested analysis would fail. In this situation the sentence is to be understood as saying ‘These gloves are good pieces of cloth for cleaning the windows’. According to Geach something is good if and only if there is an appropriate noun *N* such that that thing is a good *N*.¹⁶

One way to build Geach's idea into a semantic analysis would be to assign (7b) to (7) as its logical form, where the language of logical forms remains extensional. Here the

¹⁵ Hamann (1991), p. 666.

¹⁶ Geach endorses this view in the following passage: “Even when ‘good’ or ‘bad’ stands by itself as a predicate, and is thus grammatically predicative, some substantive has to be understood; there is no such thing as being just good or bad, there is only a being a good or bad so-and-so.” Geach (1956), p. 65.

her piano performances are terrible. That is, ‘ $\exists e [\text{dancing}(e, \text{Sue})] (\text{good}(e))$ ’ might be true and ‘ $\exists e [\text{piano-playing}(e, \text{Sue})] (\text{good}(e))$ ’ false.

But the problem of non-transparency is still with us. Imagine a world where dances are performed with a single aim: to persuade the gods to give rain for the crops. Furthermore, dancing is the only method people ever use to try to persuade the gods to give rain for the crops. In this world, not only are the dancers and the rainmakers the same people, dancing and rainmaking are the very same events.²⁰ Now, suppose Sue is a shaman-dancer-rainmaker in this world. She had perfected her dancing and she is much admired for her graceful moves. Unfortunately, she is not particularly successful in making rain: year after year, despite her best efforts there is a terrible drought. It seems to me that under such circumstances we could say that Sue is a good dancer but not a good rainmaker. But now we have a problem:

- | | |
|--|--|
| (7) Sue is a good dancer | (7c) $\exists e [\text{dancing}(e, \text{Sue})] (\text{good}(e))$ |
| <u>(11) All and only dances are rain-makings</u> | <u>(11c) $\forall e (\text{dancing}(e) \leftrightarrow \text{rainmaking}(e))$</u> |
| (12) Sue is a good rainmaker | (12c) $\exists e [\text{rainmaking}(e, \text{Sue})] (\text{good}(e))$ |

The inference on the left is invalid,²¹ but the corresponding inference run on the putative logical forms on the right is valid. So, Larson’s logical forms cannot be correct either.

There is a standard way to get around the problem of non-transparency. The idea is that the logical forms should be drawn from an *intensional language*.²² The interpretation of an intensional language distinguishes between (at least) two different semantic values of expressions: the extension and the intension. The extension of the

²⁰ One might deny that there could be a world where dancing events and rainmaking events coincide: such events could occur always simultaneously, but they would nevertheless be distinct events. I find this suggestion *ad hoc*. An anthropologist may observe dances, describe them in detail, and only later discover that their function is rainmaking. He would naturally think that he thereby found an alternative way to describe the very same events he had already observed.

²¹ I am assuming that whatever the semantics of the generic quantifier, it must validate such an inference.

expression is its ordinary semantic value (a truth-value for sentences, an object for names, a set of objects for one-place predicates, a function from sets to sets for predicate-modifiers, etc.); the intension of the expression is a function from possible worlds to its extension. One can have an intensional language where – although there are certain operators that are sensitive to the intensions of the expressions to which they are applied – quantification over possible worlds remains covert. Languages of modal logics are intensional languages of this type. Alternatively, one can have an intensional language, where reference to possible worlds is explicit. In these languages there are variables ranging over possible worlds.²³

I think intensional languages of the second type are more perspicuous, so I will use this technique to present the intensional solution to the problem posed by the transfer-inferences.²⁴ The intensionalized version of Geach’s proposal is (7d):

(7) Sue is a good dancer (7d) $(\text{good}(\text{dancer})) (\text{Sue}) (w)$

In the logical form, ‘w’ is a variable that will be assigned the possible world with respect to which the sentence is evaluated. In most cases this is the actual world, so (7d) is true iff Sue is among the individuals that are good dancers in the actual world. But (7) can also occur as a constituent in a complex sentence with modal operators, which are translated as quantifiers binding the possible world variable.

²² Cf. Montague (1970a), esp. p. 201, pp. 211 - 3, and Montague (1970b), esp. pp. 242 - 3.

²³ Cf. Lewis (1970), esp. pp. 193 - 200.

²⁴ In doing this, I don’t want to prejudge the complicated issue of ontological commitment. Philosophers tend to like non-explicit quantification over possible worlds because they believe that it fails to carry ontological commitment to those entities. I am doubtful whether pushing explicit quantification into the meta-language can help the ontologically scrupulous, but I don’t want to press the issue here. There is another typical motivation for using modal operators, rather than explicit quantification. It has been argued that natural language lacks the expressive power of full-blooded quantification over possible worlds. This claim, however, had been challenged on empirical grounds by a number of philosophers and linguists.

Such a move blocks the transfer-inference from (7) and (10) to (9). Even if the predicates ‘dancer(x)’ and ‘pianist(x)’ have the same extensions in the actual world, they have different extensions in other possible worlds. Then the intensions of the complex predicates ‘good(dancer)(x)’ and ‘good(pianist)(x)’ are different functions, and therefore the expressions ‘(good(dancer))(Sue)’ and ‘(good(pianist))(Sue)’ have different intensions too. But then there is no reason why the sentences ‘Sue is a good dancer’ and ‘Sue is a good pianist’ should have the same truth-value in the actual world.

One might wonder whether the intensional move solves the problem. What if pianists and dancers are necessarily the same individuals? Doesn’t intuition tell us that even under such circumstances Sue could be a good dancer and a bad pianist? It is hard to tell. Our intuitions might be confused here, for it is not clear whether the assumption that *necessarily* all and only pianists are dancers is conceptually coherent.²⁵ I will not pursue this line of objection for there is a more convincing one.

By assigning (7d) as logical form to (7), we implicitly commit ourselves to an explanation why the inference fails. The *reason* why Sue may be a good dancer and a pianist without being a good pianist even if all dancers are pianists and *vice versa* is that there is another possible world where someone (she, or someone else) is a dancer, but not a pianist, or a pianist, but not a dancer. This does not sound very convincing. How could the possibility that someone is a dancer, but not a pianist have anything to do with the

²⁵ But note the following objection: Arguably, ‘sell’ and ‘buy’ are expressions such that they are necessarily true of the very same events. Still, a particular transaction can be a good sell and a bad buy. So ‘The transaction between Bill and Sam was a good buy’ does not entail ‘The transaction between Bill and Sam was a good sell.’ But the argument is not fully convincing. A defender of the intensional approach may reply that selling and buying events must be distinguished. Alternatively, one might argue that

question whether Sue has to be a good pianist, given that she is a good dancer and a pianist?

It is interesting to compare this with the following case. The inference from ‘Sue is a prospective dancer’ and ‘All and only dancers are pianists’ to ‘Sue is a prospective pianist’ is also invalid. Here the intensional explanation says that the *reason* why Sue may be a prospective dancer and a pianist without being a prospective pianist even if all dancers are pianists and vice versa is that there is another possible world where someone (she, or someone else) is a dancer, but not a pianist, or a pianist, but not a dancer. In this case the explanation is much more plausible. The failure of this inference can be explained by means of possibilities, because if someone is a prospective *N*, then she *may* become, or is *likely* to become an *N*. However, similar modal consequences cannot be drawn from the claim that someone is a good *N*.²⁶

We are in a quandary. Extensional approaches to the adjective ‘good’ seem inadequate because they cannot account for its deep non-transparency. They all incorrectly predict that if the extensions of ‘*N*’ and ‘*M*’ are sufficiently closely tied, the inference from ‘*a* is a good *N*’ to ‘*a* is a good *M*’ goes through. The failure of the extensional accounts naturally suggests that we should regard ‘good’ as an intensional adjective. However, the explanation of non-transparency gained from the introduction of intensions appears to be the wrong one. The failure of the transfer inferences has nothing

although the events are the same, the actions of buying and selling are different, and that the logical forms of action verbs quantify over actions, rather than events.

²⁶ This criticism of the use of the intensional machinery in blocking certain inferences has a predecessor in the literature. Sally McConnell-Ginet writes: “But the intensional machinery does not provide a good model of how we think about *why* those walking quickly might be different from those talking quickly, even though walkers and talkers are the same. The explanation lies not in the existence of an alternative situation (where individuals have different properties), but simply in the possibility of a different sorting of the individuals, given a refinement of the sorting principles.” McConnell-Ginet (1982), p. 163.

to do with modal considerations and so, by going intensional we avoid wrong predictions only through sacrificing explanatory value. A different approach is called for.

4. Ways of Being Good

On the surface the answer to the question why someone may be a good dancer and a bad pianist is quite obvious. Dancing and playing the piano are very different skills, so there is no reason why excellence in one should have anything to do with excellence in the other. In other words, if Sue is a good dancer then she is good *at dancing*, which is perfectly compatible with her being quite bad *at playing the piano*. So, goodness – at least in these cases – does not directly attach to Sue; it attaches to her only through one or another description that is true of her. Some of these descriptions, as that she is a dancer or that she is a pianist, may specify certain roles, and she may be skillful, interesting, enthusiastic or well-paid in one of her roles, but unskillful, boring, indifferent or badly paid in another. Being good at dancing and being good at playing the piano are different *ways of being good*.²⁷ This seems to be the correct analysis of ‘good’ and many other adjectives.

As it stands, this explanation is in accordance with the analysis that regards ‘good’ as a predicate-modifier. However, theories that regard ‘good’ as a predicate-modifier tend to make a further step: they typically assume that the relevant roles are fully specified by the noun to which the adjective ‘good’ is attached. As Geach himself noted, there are cases in which this move is problematic.

²⁷ The phrase ‘way of being good’ is borrowed from Thomson (1992).

The difficulties are exemplified by sentences in which one speaks about a good event, or a good thing to happen. Geach says that such sentences often do not have a fixed content, since “ ‘event’, like ‘thing’, is too empty a word to convey either a criterion of identity or a standard of goodness.”²⁸ This means that according to Geach, the reason why ‘good event’ does not have clear truth-conditions is that ‘event’ does not have a clear content. But, surely, the content of ‘pebble’ or ‘differential-equation’ are clear enough, and still, one does not know what to make of ‘good pebble’ or ‘good differential-equation’.²⁹ The content of the sentence ‘This is a good pebble’ depends on the context of its use. In a certain situation, an utterance of this sentence might say that the pebble is good for playing marbles, in another that it is good for breaking a window.

So the noun N in a phrase ‘ a is a good N ’ is often not sufficient in determining the role in which a is supposed to be good. Is it at least always necessary? Geach certainly thinks so, since he believes that in order to interpret ‘ a is good’, one always has to provide some suitable N that will yield the standards of goodness. I think this is implausible. Consider the scenario where students are performing an experiment in the laboratory. They are trying to produce some substance that can be used in a later experiment. The teacher points to a certain blue liquid and says to the students: ‘This is good’. Is it really true that in order to interpret this, the students have to be able to come up with some noun N such that the blue liquid is a good N ? It seems that even if they know close to nothing about what the blue liquid is, and hence the nouns they could come up with do not provide standards of goodness, they can know what the teacher’s

²⁸ Geach (1956), pp. 68 - 9.

²⁹ I thank Judith Thomson for bringing this problem to my attention. Her discussion of it can be found in Thomson (1994).

utterance said. The blue liquid was said to be good for the purposes of the experiment. Of course, Geach *might* reply that the noun (or, more precisely, the nominal expression) in question is something like ‘stuff that can be used in a later experiment’. But why should we believe that the standards of goodness are provided by the content of this complex noun, rather than simply by the context in which the utterance was made?

Here is another problem with the predicate-modifier approach. Imagine a case when students are arguing about what to read next in their reading group. There are two candidates: a short book and a lengthy paper. One of the students points to the book and says: ‘This is a good book’. Another, as a reply, points to the paper and says: ‘This is good too’. What the second student said can be paraphrased roughly as ‘This is good to read too’; it *cannot* be paraphrased as ‘This is a good book too’. But standard accounts of ellipsis assume that the second sentence predicates something of the paper that is predicated of the book in the first. Given the predicate-modifier analysis of ‘good’, a suitable predicate is simply unavailable.³⁰

All of this suggests strongly that the content of the noun ‘*N*’ does not play a direct semantic role in fixing the way in which a good *N* is said to be good. I suggest that we should return to the Moorean idea that the logical form of ‘good’ is a one-place predicate of individuals. Of course, there is a price for this: one certainly cannot hold a theory according to which the extension of ‘good’ is the class of good things. There may be nothing in common to all of the good things. However, one can say that ‘good’ is an *incomplete* one-place predicate, one that is associated with a set of individuals only if

³⁰ A defender of the predicate-modifier approach will have to say that this is not a case of genuine ellipsis.

this is undoubtedly true. But if we think that the minimal aim of semantics is to provide a systematic theory of the truth-conditions of assertive uses of well-formed declarative sentences then truth-conditional effects of context are properly regarded as lying within the domain of semantic theory. This, however, does not preclude the possibility that the mechanism used to determine the value of ‘ R ’ within a context is similar (or perhaps identical) to the mechanism applied in pragmatic implications.³³

An advantage of this sort of semantics for ‘good’ is that it can account easily for the intuition that ‘Sue is a good dancer’ entails ‘There is a way in which Sue is good’ or ‘Sue is good in some respect’. On the incomplete predicate analysis, the latter sentences can be regarded as existentially quantifying over the variable within the logical form of ‘good’ whose value is normally fixed by the context. It is not clear to me how one could account for these inferences on the alternative accounts.

Another advantage is that it accounts well for cases when something is said to be a good N , but the way it is good is not determined by ‘ N ’. For example, consider a case when someone is looking for an appropriate object to prop up a television set. Upon finding a thick book, he might say ‘This is a good book.’ What his utterance says is, of course, that the book is good for propping up the television set, not that it is good *as a book*. Again, the predicate modifier approach has difficulties with capturing such intuitions.

One of the consequences of this analysis is the claim that in *some* contexts (7) could express something different from the proposition that Sue is good as a dancer and

³³ “...pragmatic factors may contribute to semantic interpretation without leading to the consequence that the idea of an independent semantic level should be given up. Pragmatic factors may trigger syntactic rules

Sue is a dancer. And indeed, one can imagine a situation in which what is at question is her moral character, not her skill at dancing. Suppose some dancers are threatening to burn down the Met, because they are unhappy about their salary. Other dancers are trying to convince the members of the first group that burning down the Met would be a bad idea. Meanwhile the singers are having a drink in a bar, and they are discussing which of their colleagues are in which group. Most of them believe that Sue is the leader of those who are threatening to set the building on fire. One of the singers tries to defend Sue, and argues that unlike most of the dancers, she is actually in the second group, trying to save the building. The singer gives several reasons for this claim, and at the end of her argument she utters the sentence ‘Sue is a good dancer’, perhaps putting a strong emphasis on ‘good’.³⁴

But even if there are intricate examples where the proposition that Sue is a good dancer does not say that she is good at dancing, in normal contexts it does: ‘good’ within ‘good *N*’ is normally has the content of ‘good as an *N*’. This means that in order to accept (7e) as the logical form of (7), one has to slightly expand the notion of context informally introduced at the beginning of this paper. There I said that the content of previous utterances might play a role in determining the content of an expression. But this is not enough. Now I suggest that the content of certain expressions within the utterance may be part of the context that contributes to determining the content of other

as well and in spite of this we would not like to say that there is no such thing as syntax and everything should be incorporated into pragmatics.” Kiefer (1978), p. 161.

³⁴ One might object that in this example I exploited an ambiguity. It seems plausible that ‘good’ has a special sense in which it means something like ‘morally good’, and it is this ambiguity, rather than some sort of context-dependence that properly accounts for the case. But I don’t think it is right to say that in my example the defender of Sue declared her to be morally good. This is revealed by the fact that uttering ‘Although morally reprehensible, Sue is a good dancer’ (again, giving a strong emphasis to ‘good’) would not have been a contradictory assertion in the given context.

expressions within the *same* utterance. The content of ‘good’ in an utterance of (7) may depend, in part, on the content of ‘dancer’.³⁵

5. Varieties of Incompleteness

The story about the semantics of ‘good’ told in the previous two sections has obvious implications for Travis’s example concerning the color of the painted leaf. If it can be made plausible that the analysis given for ‘good’ – according to which its logical form is a contextually incomplete predicate – applies to other adjectives, including ‘green’, then we may have a semantic account of these expressions that does not violate the context thesis. If ‘green’ is an incomplete predicate, the context-dependence of ‘The leaf is green’ is attributable to the context-dependence of ‘green’.

The resistance to the predicative analysis in the semantic literature is largely based on the non-transparency of most adjectives. Once it is acknowledged that the invalidity of these inferences can be attributed to contextual incompleteness, the resistance loses much of its force. The advantages of a uniform predicative approach to all adjectives are considerable. First of all it is uniform: it assigns the same kind of semantic value to all adjectives.³⁶ This is a theoretical gain for it keeps the syntactic and semantic structures close to one another. Second, assuming a broadly Davidsonian semantics for adverbs, we can account for similarities between adjectival and adverbial

³⁵ Fortunately, there do not seem to be cases when the content of a noun ‘*N*’ within a ‘good *N*’ depends on the content of ‘good’ in that phrase. This accounts for the strong intuition that within a noun phrase of the form [_{NP} *AN*], the semantic status of the adjective and the noun are different. The content of the adjective is influenced by the content of the noun, but not the other way around.

³⁶ There are many who defend a double classification for adjectives; cf. Parsons (1972) and Siegel (1976).

modification. Adjectives and adverbs can both be treated as predicates: the former as predicates of individuals, the latter as predicates of eventualities.³⁷ Proponents of a unified predicate-modifier approach have no similarly straightforward account to offer about modification in general.³⁸ Finally, there is some psycholinguistic evidence that the predicative use of adjectives is acquired (and, in the case of aphasia- and dysphasia-patients, re-acquired) first.³⁹ This suggests that the predicative use of adjectives is more basic.

Nevertheless, the predicative analysis cannot be universally applied to all adjectives. Some (e.g. ‘utter’, ‘former’, ‘chief’, ‘final’, etc.) have no predicative occurrences, so the assumption that they stand for incomplete predicates is extremely implausible. And ‘Sue is a retired pianist’ does not have the same content as ‘Sue is retired (as a pianist) and Sue is a pianist’.⁴⁰ I suggest that we set the adjectives which do not support the inference pattern from ‘*a* is *AN*’ to ‘*a* is *N*’ aside, and accept for all the others an incomplete predicate analysis. In this way we maintain a uniform semantic analysis for the vast majority of adjectives.⁴¹

³⁷ Eventualities are events, states and processes.

³⁸ Classic proponents of the uniform predicate-modifier approach include Montague (1970a) and Cresswell (1976).

³⁹ Cf. Osgood (1971).

⁴⁰ The fact that ‘retired’ has predicative uses is something of an anomaly. Note that ‘Sue is a retired pianist’ is hardly distinguishable from ‘Sue is a former pianist’, even though the former entails ‘Sue is retired (as a pianist)’ but the latter does not entail *‘Sue is former (as a pianist)’. The difference between ‘retired’ and ‘former’ is probably due to the fact that ‘retired’ is derived from a participle, whereas ‘former’ isn’t.

⁴¹ Is it legitimate to set these cases aside? The trouble, of course, is that it might turn out that a semantic account which covers these cases can account for those dealt with the predicative analysis as well. Still, I think it would be bad methodology to assimilate all adjectives to the relatively few which fail to support the inference from ‘*a* is *AN*’ to ‘*a* is *N*’. For then the fact that these inferences are valid for the vast majority of adjectives would have to be explained via *ad hoc* meaning-postulates. Those who insist on treating all adjectives as they would treat fairly exceptional cases are “generalizing to the worst case”, just as Montague did when he assigned intensional types to all expressions.

Contextual incompleteness in adjectives has different dimensions. The interpretation of *scalar* adjectives, like ‘tall’, ‘heavy’, ‘fast’, ‘expensive’ or ‘old’, requires a comparison class. Short basketball players are usually tall people, light whales are heavy animals and – for many of us – cheap airplane tickets to Europe are expensive gifts. The interpretation of *evaluative* adjectives, like ‘lucky’, ‘delicious’, ‘time-consuming’, ‘simple’ or ‘fitting’, requires that context provide an individual or group of individuals from whose perspective the evaluation is made. A particular event can be fortuitous for me, but not for others; an exercise can be simple for the expert and taxing for the novice, the outcome of an election may be lucky for one party, but not for the country. And of course, there are adjectives that belong to both of these categories; ‘expensive’ and ‘time-consuming’ would be good examples.

I suggest that different dimensions of incompleteness correspond to different sorts of variables in the logical form. The variables can receive their values from the context but also from linguistic material. In English, prepositional phrases tend to perform the latter job. For example, ‘clever *at* doing cross-word puzzles’ and ‘time-consuming *as* a short-paper topic *for* graduate students’ are wholly specified one-place predicates. Accordingly, I suggest that the logical form of ‘clever’ is ‘(clever(R))(x)’, the logical form of ‘time-consuming’ is ‘(time-consuming(C₁, C₂))(x)’, where ‘R’ is a role-variable standing for a way of being clever, ‘C₁’ is a class-variable, standing for a class of with respect to which something counts as time-consuming, and ‘C₂’ is another

class-variable, standing for a class from whose perspective something counts a time-consuming.⁴²

For predicates that are fully specified, transfer-inferences are valid. For example, from ‘Sue is a clever dancer’ and ‘Sue is a pianist’ it *does* follow that ‘Sue is a clever pianist’, *provided* that e.g. the context had already specified that ‘clever’ is to be understood as ‘clever at solving crossword-puzzles’ in both the first premise and the conclusion.⁴³ Of course, if only cleverness in solving crossword puzzles is concerned, it would be quite misleading to say that Sue is a clever dancer or a clever pianist. One should put it rather as ‘Sue, who is a dancer (pianist) is clever’. The reason for this is that the sentence ‘Sue is a clever dancer’ strongly *suggests* that the kind of cleverness that is ascribed to Sue has something to do with the fact that she is a dancer.

So what should we say about ‘green’? It is certainly not an evaluative adjective: no matter how much one believes in the possibility of an inverted spectrum, it is incorrect to say that a leaf is green *for* me. We use color terms as if color were an objective, observer-independent feature of the world. We have another construction, namely ‘seems green’ which is correctly used in cases when the speaker wants to express uncertainty with respect to veridicality of her subjective experience.

⁴² It is important to notice that the prepositional phrase that can complete adjectives can modify different constituents of a sentence and that the results are not all equivalent. ‘This mouse, for a mouse, is big’ is equivalent to ‘This mouse is big for a mouse’, and ‘This, for a mouse, is big’ is equivalent to ‘This is big for a mouse’, but the first two sentences are not equivalent to the last two. The first two sentences entail that the object demonstrated is a mouse, the last two do not. Cf. Bartsch (1987), p. 16.

⁴³ Bartsch (1987) notices that ‘John’s paper is stylistically good’ entails ‘John’s paper is good’ if we evaluate the premise and the conclusion within a context in which the same respects of goodness are relevant. (p. 5) Nevertheless, instead of the approach advocated here, she endorses a context-sensitive predicate-modifier account of the semantics of adjectives.

‘Green’ is probably a scalar adjective.⁴⁴ But setting degrees of greenness aside, one might be tempted to think that ‘green’ is a paradigm example of a contextually complete adjective. To be green is one of the most obvious examples of a feature that is *simple*. It is often difficult to decide whether something is green because of the condition of the object, the environment, or the nature of our perceptual mechanism; there are also borderline cases, when there might be no fact of the matter whether something is green. But all things that are green are green in the very same way.

I think this intuition is misleading. There are at least two ways in which an apple can be green: from the outside, or from the inside. In the former case, it can be ripe, in the latter it cannot.⁴⁵ There are at least three ways in which a book can be green: due to having a green dust jacket, a green cover, or green pages. And a corridor can be green in many ways: having green walls, or green ceiling, or green carpet, or green doors, etc.⁴⁶

An object is green if some contextually specifiable (and presumably sufficiently large) *part* of it is green. The logical form of ‘green’ is ‘(green(C, P)) (x)’, where ‘C’ is a class standing for a comparison class and ‘P’ is a variable standing for a certain part of the object.⁴⁷ It seems to me that the case of the painted maple leaf fits this pattern. If one is sorting leaves for decoration, what matters is the color of the outside, if one is

⁴⁴ Languages may vary in how they categorize color adjectives, there might be even variation within the same language. For example, in Hungarian comparative forms of basic color adjectives are all acceptable, but comparative forms of complex color adjectives are not. Cf. Kiefer (1978), p. 154.

⁴⁵ Could it be that in one of these cases the apple is not literally green? This is unlikely, for there seems to be non non-arbitrary way to decide *which* of the two cases are cases on non-literal greenness.

⁴⁶ Lahav (1989) rejects the idea that the context sensitivity of ‘red house’ can be accounted for in this manner. He points out that we might be interested in the color of the inside of a house as much as we are interested in the color of its external surface, the color of the rooms does not normally count towards the color of the house. Although this observation is correct, it does not undermine the analysis. In some contexts it may well be correct to call a house red on account of its interior. The fact that such contexts are unusual is of no semantic concern.

⁴⁷ In this regard, ‘bald’ behaves much like color adjectives. One can be bald on the top of one’s head, in the front, in the back, none of which means that one is bald altogether.

trying to identify the species, what matters is what we find under the camouflage. This suggests that the context-dependency that appears in Travis's example is a relatively easily characterizable kind: it is a matter of different contextually specified values for the variable 'P'.

6. Conclusion

My strategy of defending the compositionality of content against Travis' challenge was to locate the source of the context-dependence of 'The leaf is green' in the adjective 'green'. I argued that the best semantics for 'green' postulates a variable within the logical form of this word. Since I also argued that – despite Travis' contention – this challenge couldn't be broadly generalized beyond sentences containing adjectives, the answer is to the point.

One might wonder whether in responding to similar challenges we have to postulate variables in other lexical items as well. I think we do. Consider, for example, the sentence 'Every leaf is green'. Arguably, many different things can be said with this sentence, depending on what the contextually relevant domain of quantification is supposed to be. And of course, one may have the *prima facie* intuition that none of the lexical items has a context-dependent content. Here again, I think the first intuition is misleading; a general account of domain restriction may demand that we postulate the presence of a domain variable within the logical form of 'leaf'.⁴⁸

⁴⁸ We argue for this view in Stanley and Szabó (2000).

So, the strategy of defense of compositionality pursued here is likely to lead to a lexicon where many, perhaps most entries contain contextual variables. Whether the defense of compositionality I offered is ultimately convincing depends largely on the plausibility of such a lexicon. And that, in turn, depends on whether we can provide an account of how the values of these variables are to be determined in context. At the beginning of the paper I set foundational questions regarding the nature of context and our knowledge of its features aside. At the end, I concede that as long as these issues remain unresolved, the semantic proposal I gave remains questionable. But even if that is so, I hope that I succeeded in pointing out where the real issues lie regarding the compositional interpretation of adjectives.

References

- Bartsch, Renate. (1987) 'Context-dependent Interpretations of Lexical Items' in J. Groenendijk, D. de Jongh and M. Stokhof eds., *Foundations of Pragmatics and Lexical Semantics*, 1 - 26. Dordrecht: Foris Publications.
- Cresswell, Max. (1976) 'The Semantics of Degree' in B. Partee ed., *Montague Grammar*, 261 - 92. New York: Academic Press.
- Geach, Peter. (1956) 'Good and Evil' reprinted in P. Foot ed., *Theories of Ethics*, 64 – 74. Oxford: Oxford University Press.
- Hamann, C. (1991) 'Adjectival Semantics' in A. von Stechow ed., *Handbuch der Semantik*. Berlin: Walter de Gruyter.
- Kaplan, David. (1977) 'Demonstratives. An Essay on the Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals' reprinted in J. Perry and H. Wettstein eds., *Themes from Kaplan*, 1989, 481 – 563. Oxford: Oxford University Press.
- Kiefer, Ferenc. (1978) 'Adjectives and Presuppositions' *Theoretical Linguistics* **2**: 135 – 73.
- Lahav, Ron. (1989) 'Against Compositionality: The Case of Adjectives' *Philosophical Studies* **57**: 261 – 79.
- Larson, Richard. (1998) 'Events and Modification in Nominals' in D. Strolovitch and A. Lawson eds., *Proceedings From Semantics and Linguistic Theory (SALT) VIII*, 145 – 168. Ithaca, NY: Cornell University Press.
- Lewis, David. (1970) 'General Semantics' reprinted in *Philosophical Papers, vol. 1*, 1983, 189 – 229. Oxford: Oxford University Press.
- McConnell-Ginet, Sally. (1982) 'Adverbs and Logical Form: A Linguistically Realistic Theory' *Language* **58**: 144 - 84.
- Montague, Richard. (1968) 'Pragmatics' reprinted in R. Thomason ed., *Formal Philosophy*, 1974, 95 – 118. New Haven: Yale University Press.
- Montague, Richard. (1970a) 'English as a Formal Language' reprinted in R. Thomason ed., *Formal Philosophy*, 1974, 188 – 221. New Haven: Yale University Press.
- Montague, Richard. (1970b) 'Universal Grammar' reprinted in R. Thomason ed., *Formal Philosophy*, 1974, 222 – 46. New Haven: Yale University Press.

- Moore, G. E. (1903) *Principia Ethica*. Cambridge: Cambridge University Press.
- Osgood, Charles, E. (1971) 'Where do Sentences Come From' in D. Steinberg and L. Jacobovits eds., *Semantics: An Interdisciplinary Reader in Philosophy, Linguistics and Psychology*, 497 - 529. Cambridge: Cambridge University Press.
- Parsons, Terence. (1972) 'Some Problems Concerning the Logic of Grammatical Modifiers' in D. Davidson and G. Harman eds., *Semantics of Natural Language*, 127 – 41. Dordrecht: Reidel.
- Siegel, Muffy E. A. (1980) *Capturing the Adjective*. New York: Garland Press.
- Stanley, Jason and Zoltán Gendler Szabó (2000) 'On Quantifier Domain Restriction' *Mind and Language* **15**: 219 - 61.
- Stalnaker, Robert. (1970) 'Pragmatics' reprinted in *Context and Content: Essays on Intentionality in Speech and Thought*, 31 - 46. Oxford: Oxford University Press.
- Szabó, Gendler Zoltán (2000) 'Compositionality as Supervenience' *Linguistics and Philosophy* **23**: 475-505.
- Thomson, Judith Jarvis (1992) 'On Some Ways in Which a Thing Can Be Good' *Social Philosophy and Policy* **9**: 96 – 117.
- Thomson, Judith Jarvis (1994) 'Goodness and Utilitarianism' in *Proceedings of the American Philosophical Association* **67**: 7 - 22.
- Travis, Charles (1994) 'On Constraints of Generality' in *Proceedings of the Aristotelian Society*. New Series **44**: 165 - 88. London: Aristotelian Society Publications.