

Color, context, and compositionality

Christopher Kennedy · Louise McNally

Received: 18 July 2008 / Accepted: 15 October 2008
© Springer Science+Business Media B.V. 2009

Abstract Color adjectives have played a central role in work on language typology and variation, but there has been relatively little investigation of their meanings by researchers in formal semantics. This is surprising given the fact that color terms have been at the center of debates in the philosophy of language over foundational questions, in particular whether the idea of a compositional, truth-conditional theory of natural language semantics is even coherent. The challenge presented by color terms is articulated in detail in the work of Charles Travis. Travis argues that structurally isomorphic sentences containing color adjectives can shift truth value from context to context depending on how they are used and in the absence of effects of vagueness or ambiguity/polysemy, and concludes that a deterministic mapping from structures to truth conditions is impossible. The goal of this paper is to provide a linguistic perspective on this issue, which we believe defuses Travis' challenge. We provide empirical arguments that color adjectives are in fact ambiguous between gradable and nongradable interpretations, and that this simple ambiguity, together with independently motivated options concerning scalar dimension within the gradable reading accounts for the Travis facts in a simpler, more constrained, and thus ultimately more successful fashion than recent contextualist analyses such as those in Szabó (Perspectives on semantics, pragmatics and discourse: A festschrift for Ferenc Kiefer, 2001) or Rothschild and Segal (Mind Lang, 2009).

Keywords Color terms · Compositionality · Indexicality · Contextualism · Ambiguity · Gradability · Classifying properties

C. Kennedy (✉)
Department of Linguistics, University of Chicago, 1010 E 59th Street, Chicago, IL 60637, USA
e-mail: ck@uchicago.edu

L. McNally
Departament de Traducció i Ciències del Llenguatge, Universitat Pompeu Fabra,
c/Roc Boronat, 138, 08018, Barcelona, Spain

1 Green leaves

The focus of this paper is the semantics of color adjectives and the implications of examples like (1) for the view, taken for granted by most work in linguistic semantics, that the semantic content of a sentence is compositionally derived from the meanings of its components and determines a set of truth conditions.

(1) The leaves are green.

In a series of influential papers, Charles Travis has brought to light a particular form of truth conditional variability associated with color terms that appears to present a challenge for this view (see e.g., [Travis 1985, 1994, 1997](#)).¹ The following passage illustrates the phenomenon:

A story. Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, 'That's better. The leaves are green now.' She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those.' But now Pia speaks falsehood. ([Travis 1997](#), p. 89)

This scenario appears to show that distinct utterances of the words in (1), said in order to describe the same scenario (the relation between the leaves and a particular color), can be associated with distinct truth values. On the assumption that (1) has a single lexico-syntactic analysis—that its surface form does not hide an underlying structural or lexical ambiguity²—this fact calls into question the hypothesis in (2).

(2) The semantic value of a sentence is a function from facts about the world to a unique value in $\{0,1\}$, the exact nature of which is determined by its syntax together with the semantic values of its constituents.

Travis' conclusion is that this hypothesis is incorrect: even on a "stipulated semantics", the sentence in (1) "is compatible with various distinct conditions for its truth" (*ibid.*).

¹ The case of color adjectives is just one type of example that Travis invokes to make essentially the same point; other examples that figure prominently in his arguments involve adjectives like *round* in (ia) (said of a squash ball traveling at full speed as it meets a wall) and existential statements like (ib) (said of a refrigerator containing nothing but a small puddle of milk on the bottom).

- (i) a. That squash ball is round.
b. There is milk in the refrigerator.

We focus on color terms in this paper both because they have received relatively little attention in the formal semantics literature, and are therefore deserving of scrutiny, and because they seem to us more problematic than cases like those in (i). The truth/falsity of (ia), for example, is quite plausibly due to the fact that *is round* can have either a stage- or an individual-level interpretation; these interpretations arguably involve distinct semantic compositions (see [Carlson 1977](#); [Chierchia 1995](#); [Kratzer 1995](#) for various analyses). (ib), on the other hand, is on our judgment quite simply true of any scenario that involves a quantity of milk in the refrigerator, with a Gricean analysis accounting for any apparent evidence to the contrary (such as Travis' example of Odile's apparently false utterance of (ib) in reference to the milk puddle in the fridge, as a response to his friend Hugo's request for milk to put in his coffee; cf. [Sainsbury 2001](#)).

² For the purposes of this discussion we include under structural ambiguity the compositional ambiguity that arises when a sentence contains more than one quantificational expression.

According to him, the semantic value of a sentence at most imposes some necessary conditions under which it may be true (as well as conditions under which it may be used), but those conditions need not be sufficient, and the content of the sentence does not define a function from contexts to truth.

This is a significant challenge, and deserves to be taken quite seriously, as (2) is a foundational assumption of linguistic semantics. One response is to simply deny the judgments and instead to take the position that the two utterances of (1) have the same truth conditions in Travis' story, and so the same truth value, contrary to initial appearances. This is the position advocated by [Sainsbury \(2001\)](#), who argues that the content of the adjective *green* is nonspecific enough to render (1) true in both uses, and that our intuitions about the falsity of (1) when said to the botanist are due to a misguided tendency to assume that it is made true in a particular way. While we accept that Sainsbury's response may be correct for some cases (such as the milk example in note 1), we, along with all the native speakers we have consulted, find it very difficult to deny Travis' empirical claim that (1) is false as a response to the botanist. We will therefore proceed on the assumption that denying the judgments is not an option.

This leaves us with two options to pursue in the defense of (2). One is to derive truth conditional variability in (1) via the normal mechanisms of a compositional, truth conditional semantics by hypothesizing that one of the constituents out of which (1) is constructed has a context dependent denotation; this is the approach taken by [Szabó \(2001\)](#), who posits a hidden variable in the denotation of color adjectives, and by [Rothschild and Segal \(2009\)](#), who treat color adjectives as full-blown indexical predicates. The other is to show that, contrary to initial assumptions, the superficial form of (1) does in fact hide an underlying ambiguity, and that this ambiguity accounts for the different judgments in Travis' example. We know of no response to Travis in the literature that takes this second approach; our purpose in this paper is to show that it is the correct one for the particular examples Travis offers (though we will also claim that the use of hidden variables provides a good account of certain other aspects of the variability in color term interpretation).

We base our arguments on a careful examination of a broader array of linguistic data than has previously been considered in the discussion on color terms. By teasing apart those aspects of the color terms' meaning that are vague or context dependent from those which involve true ambiguity, and by giving these components separate treatments, we provide an overall account of the semantics of color terms which is more constrained, and situates these terms more successfully within the overall picture of the lexical semantics of adjectives, than do unadorned contextualist approaches.³

The structure of the paper is as follows. In the next section, we outline the analyses of color adjectives proposed by [Szabó \(2001\)](#) and [Rothschild and Segal \(2009\)](#), showing that while they both succeed in preserving a compositional, truth conditional semantics in the face of Travis' example, they face empirical challenges that suggest

³ There is a third option, advocated in [Predelli \(2005\)](#), which is to say that the content of a sentence containing a color adjective determines a fixed set of truth conditions, but to allow for the possibility that the truth value of the sentence could still vary according to different points of evaluation. While this may be an option that we want to consider in general, since the point of this paper is to argue that Travis' examples involve a true ambiguity, and to explain the nature of this ambiguity, we will not evaluate it here.

that neither is the full answer. In Sect. 3, we introduce new data from the interaction of color adjectives and comparative/degree morphology which raise problems for the pure indexical accounts and provide compelling reasons to believe that color adjectives have both gradable and nongradable meanings, and that this distinction correlates with the truth conditional variability observed by Travis. In Sect. 4 we develop a semantic analysis of color adjectives on which the gradable reading corresponds to the simple property of manifesting color, while the nongradable reading corresponds to the property of manifesting, at some point or other, color which is nonaccidentally correlated with some other relevant property of the object to which the color is ascribed. Finally, in Sect. 5, we consider the broader implications of our proposals.

2 The indexical response(s)

As noted above, one response to Travis' arguments is to show that the property expressed by the color adjective *green* is itself context dependent in just the way we need to get the truth conditional variability observed for (1). We discuss two variants of this approach which differ in where they locate context dependence.

2.1 Hidden variables

On the first view, advocated by Szabó (2001), a color adjective is analyzed as containing a parameter whose value can be fixed in different ways in different contexts. Building on the analysis of adjectives like *skillful* and *talented*, which might plausibly be analyzed as having a parameter (beyond the standardly assumed "comparison class" variable for gradable/vague predicates) that specifies the way in which an object is considered skillful or talented (made explicit with an *as*-phrase: *skillful/talented as a violinist vs. as a linguist*), Szabó proposes that color adjectives have denotations as in (3), where C is a comparison class and P is a variable that picks out the part of x that the property represented by **green** is applied to in order to assess truth.⁴

$$(3) T(\textit{green}) = \lambda P \lambda C \lambda x. \mathbf{green}(P)(C)(x)$$

The value of P is usually fixed by the context, so *is green* is true of an object a just in case a contextually determined (and presumably sufficiently large) part of a is green.

According to Szabó, this analysis can handle Travis' example by simply setting the value for P in the right way. When Pia first uses (1) to describe her leaves, she fixes the value of P in such a way that it picks out the surface area of an object; *is green* thus denotes the property of being superficially green, and so is true of the painted leaves. On the other hand, assuming Pia has correctly understood the intentions of her botanist friend, her second utterance of (1) should involve fixing the value of P in such a way that it picks out the entirety of the objects to which the predicate applies; *is green* then denotes the property of being entirely green, and is now false of the (merely) painted leaves.

⁴ $T(\alpha)$ stands for the translation of α into the higher order predicate calculus we use here to represent denotations.

While we share Szabó's intuition that the part structure of the object to which a color term applies is a relevant parameter in the interpretation of these terms, we feel that the proposal in (3) is insufficient both as a response to Travis and as a fully comprehensive account of the semantic properties of color adjectives. We will put off a demonstration of the second point until Sect. 3, where we show that the interpretation of particular color adjectives can vary systematically in ways that cannot be captured merely by reference to the part structure of their arguments (or to comparison classes). The first point is shown by the fact that the judgments about (1) remain the same if, instead of merely painting her leaves, Pia immerses them in a dye that has the effect of rendering them green throughout, or at least green in all the parts of that are green in naturally green leaves. The botanist would still be justified in calling Pia a liar if she discovered that the leaves Pia had handed her while uttering (1) turned out to have been treated in this way. Szabó's semantic analysis, which links truth conditional variability of color terms to differences in which parts of an object manifests the color, cannot explain this result.

2.2 Indexical predicates

On the second view, developed in most detail by [Rothschild and Segal \(2009\)](#), color adjectives are treated as full-blown indexical predicates that denote distinct properties in different contexts of utterance. On this view, the meaning of the adjective *green* is represented as the indexed property in (4a), which has the semantics in (4b) (our notation differs in trivial details from Rothschild and Segal's).

- (4) a. $T(\textit{green}) = \mathbf{green}_i$
 b. For any x , $\mathbf{green}_i(x)$ is true in a context C_i iff x is green according to the standard for greenness in C_i

Rothschild and Segal's notion of context is rather technical: The painter's use of green would correspond to one context in their sense, while the botanist's would correspond to another, even though both uses of the color term might be found in the same conversation. Thus, the representation of (1) as uttered by Pia to refer to her painted leaves might be as in (5a), while the representation of the same sentence as interpreted by the botanist might be as in (5b) (any combination of indices being in principle possible, as long as they are not identical in the two cases):

- (5) a. $\mathbf{green}_i(\textit{the leaves})$
 b. $\mathbf{green}_j(\textit{the leaves})$

One clear advantage of this kind of analysis is that it provides for the kind of context dependence we need to handle the Travis examples without having to posit hidden variables.⁵ However, this simplification of the denotation of a color adjective comes at a significant cost: it is not clear what constraints there are on possible valuations of

⁵ Indeed, much of Rothschild and Segal's paper is devoted to showing that there is little positive evidence for implicit variables of the sort assumed by Szabó in the logical form of sentences with color adjectives. Since our proposal is largely independent of the question of whether color adjectives have denotations that include implicit variables, we will mostly ignore these questions here.

a particular adjective in different contexts. Rothschild and Segal are aware of this, but feel that the question of what the valuation of an adjective like *green* is in a particular context is a “matter of psychology” rather than something that semantic theory should explain.

One initial argument that psychology is not all that is involved emerges when we consider whether all contextually and psychologically plausible valuations for a color adjective are in fact attested. One such valuation (perhaps of interest to the color theorist) could be “is *color* in any of the senses that the term is ever used”. Using *red* as our example, let us call this denotation \mathbf{red}_g . (In fact, Rothschild and Segal suggest that such an interpretation exists as a context-independent interpretation that subsumes the context dependent interpretations.) Such a property is coherent, and can be used to characterize objects that are red in different ways.

Consider, for example, a context involving two traffic signals, A and B. The uppermost light of A is illuminated and glowing red (indicating “stop”), but the body of the signal is painted black. In contrast, the lowermost light of B is illuminated and glowing green, and the body of the signal is painted red. This situation could be felicitously described using the sentence in (6).

- (6) Traffic signal A is red in one of the senses that the term is used, and so is traffic signal B.

The possibility of *so*-anaphora in this example is important. This is a species of property anaphora that requires identity of sense between antecedent and anaphor (Lakoff 1970; Zwicky and Sadock 1975), so the fact that (6) is acceptable indicates that the same property is being attributed to the two lights, namely the property \mathbf{red}_g , expressed periphrastically. But if this property were also a possible denotation of unmodified *red*, then it should be just as easy to interpret (7) in the same way as (6), contrary to fact: here *red* is either true in the first conjunct and false in the second, or vice-versa, depending on how it is interpreted.

- (7) Traffic signal A is red, and so is traffic signal B.

Thus, while there would appear to be nothing conceptually wrong with a property such as \mathbf{red}_g , natural language evidently does not lexicalize this sort of meaning, with or without the help of indexical parameters.

Perhaps even more strikingly problematic for a pure indexical analysis is the impossibility of understanding \mathbf{green}_i as meaning exclusively “painted green”. Consider (8), uttered in a context in which leaf A and leaf B are objectively the same rich shades of green, but leaf A has been painted to look that way, while leaf B is naturally that color.

- (8) Leaf A is \mathbf{green}_i but leaf B isn't.

The second conjunct in (8) involves ellipsis, which like *so*-anaphora requires identity of sense (see Grinder and Postal 1971; Hankamer and Sag 1976, in addition to the references cited above). This means that the very same property used in the first conjunct must be “copied in” as the value of the missing predicate in the second conjunct. If \mathbf{green}_i in the first conjunct can be understood to mean “painted green”, as predicted by the indexical analysis, then (8) should have a reading that is equivalent to (9).

(9) Leaf A is painted green, but leaf B isn't.

(9) is both true and informative in the context under consideration (it's a way of saying which leaf has been painted and which one hasn't). Our intuitions about (8), however, are that it is quite clearly false in the described context: if we accept the truth of the first conjunct, then given the leaves' objective color identity, we must reject the negative attribution of the same property to the naturally green leaf.⁶

The situation is of course different if instead leaf A is the naturally green one and leaf B is the painted one: this is just the Travis context, and (8) can be used to truthfully describe this situation. At the same time, one could respond to an utterance of (8) in this context by saying "no, that's not true, look at them, they're both green", and by that mean that they both have the same objective color properties (independent of how they got that way). There is thus a real ambiguity in this example, but of a more limited sort than the indexical account would lead us to expect: *green* can denote a property that is true of naturally green leaves and false of painted ones, or a property that is true of leaves with certain objective color features (independent of how they got that way), but not a property that is true only of objects that got those features in virtue of being painted (or dyed or whatever).

What these facts show is that while there is indeed some flexibility in the interpretation of color adjectives, the range of their truth-conditionally distinct interpretations (or their distinct contributions to the truth conditions) is not unlimited, and certainly not as wide-ranging as a bare-bones indexical analysis—not to mention Travis-style contextualism—would lead us to expect. An indexical account could in principle be constrained to deal with the facts by supplementing it with an independently motivated theory of what counts as a psychologically plausible valuation for an indexical predicate and what does not. However, in the absence of such a theory, the observed constraints on the interpretation of color adjectives suggest a different line of investigation, one which explores the possibility that certain parameters of color term meanings may be fixed after all. In order to show that this is in fact the case, we will look at a broader array of data than have previously been considered and demonstrate that there are correlations between the valuations a color adjective takes on and its linguistic behavior. These correlations involve the (un)acceptability of modification by degree morphology, as well as the surprising fact that color terms can sometimes be applicable even when the described object does not manifest the color in question at all. We will use these facts to argue that color adjectives are in fact ambiguous, contrary to

⁶ Daniel Rothschild (personal communication) suggests that (8) could be judged true in a situation in which leaf A is the painted one if being painted green (or not) is somehow an important feature for classification or categorization (relative to some interests). For example, imagine an artist who has been painting grey paper leaves so that they exactly match a pile of naturally green leaves. A burst of wind mixes her piles, and she needs to sort through them to separate the painted leaves from the real ones. According to Rothschild, in this context she could hold up a painted leaf A and a real leaf B and utter (8) as a way of telling her assistant how to categorize them. We are not in full agreement with Rothschild's judgments, but the fact that his example crucially involves a notion of classification is significant given the analysis we will present in the next section, as it allows for the possibility of treating this use of *green* as a special case of the nongradable semantics for *green* to be introduced below. If this is correct, it is still consistent with our general position that the range of interpretations of color adjectives is more constrained than an indexical account would lead us to expect.

the assumptions of the analyses discussed above, in exactly the way that we need to respond to Travis' challenge.

3 Color adjectives are ambiguous

3.1 Degree of color

We have seen that the two analyses of color adjectives as context dependent terms presented in the previous section provide plausible ways of preserving a truth conditional semantics in the face of Travis' arguments, though neither is without its own internal flaws. Beyond these worries, however, both accounts consider a fairly limited set of examples involving color terms, focusing almost exclusively on examples that have the grammatical structure of (1): unmodified color adjectives in predicate position.⁷ The same is true of Travis' work on this issue. This raises the question of whether a more extensive examination of the linguistic properties of color adjectives might shed new light on the puzzle we are faced with. In fact, we do not have to go far to find a context that distinguishes different uses of color adjectives in just the way we want: the interaction with comparative and degree morphology.

Let us consider a modified version of the story of Pia and her leaves. Now she has a pile of painted leaves of varying shades of green (pile A) as well as a pile of naturally green leaves, also of varying shades (pile B). Pia's artist friend walks in and asks if she can have some green leaves for a project. Pia invites her to sort through the piles and take whichever leaves she wants. In sorting through the piles, the artist might utter any of the sentences in (10) in reference either to leaves from pile A or to leaves from pile B, as appropriate based on the way that they manifest green: the particular combination of hue, saturation, and brightness, extent of color, and so forth (more on this below).

- (10) a. These leaves are green.
 b. These leaves are greener than those.
 c. These leaves aren't as green as those.
 d. These leaves are less green than those.
 e. These leaves are not green enough.
 f. These leaves are too green.

⁷ Occasionally in these works the color adjectives appear as prenominal modifiers (e.g. *the green leaves* vs. *the leaf is green*). However, such examples are not the best test cases for claims concerning compositionality, as some color terms form compounds with the nouns they modify. Consider (ia), for example, which has a reading that cannot be paraphrased as (ib).

- (i) a. I want green peppers for this recipe.
 b. I want peppers that are green for this recipe.

The contrast between these two examples suggests that *green pepper* has a reading on which it is a true compound with a not entirely compositional semantics, and that it has the status of a lexical item in English as the name for a certain type of pepper. As it is not always so easy to exclude the possibility that a given color term forms a lexicalized compound with the noun it modifies, we focus on predicative uses, where the attested truth conditional variability indicates that something needs to be said about the adjective independently of whatever interpretive effects might arise in the interpretation of compounds.

- g. These leaves are completely green.
- h. These leaves are perfectly green.
- i. These leaves are pretty/really green.
- j. These leaves are not so green.

What is important to observe is that for the artist who is interested in the colors of the leaves in her composition, any of these sentences would in principle be felicitous. Furthermore, (10a) is true of all of the leaves—both the painted ones and the natural ones—provided they are ‘green enough’. The only issue is how green they are, or maybe how much of each of them is green; why they are green (i.e. because they are naturally or artificially so) is irrelevant.

The situation is different for the botanist. She is perfectly justified in continuing to reject (the words in) (10a) as a false description of the painted leaves, while accepting it as true of the natural leaves. However, if these are her judgments about (10a), then none of the examples in (10b–i) are acceptable as descriptions of any of the leaves. That is, she cannot point to pile B (the naturally green leaves) and utter (10a) with the intended meaning (that the leaves are naturally green), and then strengthen or reiterate her point by pointing to pile A and uttering (10e) or (10j). Similarly, there is no way for her to use (10b) to justify her selection of the naturally green leaves over the painted ones, or (10c–d) to justify rejection of the latter, strictly on the basis of their biological properties. In short, once she starts making comparisons or comments on degree, the painted/natural distinction is out of the picture; all that is relevant are relative degrees of some objective manifestation of color.⁸

What these facts show is that there is a semantic difference between occurrences of *green* that are used to distinguish between objects on the basis of why they are green (e.g., chlorophyll vs. paint) and instances that are used to distinguish between objects on the basis of how they are green (depth of hue, proximity to a prototype, extent of color, etc.). Each of (10b–j) involves the combination of the color adjective with a different element from the set of English degree morphemes, all of which require the adjective they combine with to be gradable. The fact that (10b–j) are acceptable when (10a) is true of both sorts of leaves shows that on this use, the color term is gradable; the fact that (10b–j) are unacceptable when (10a) is true only of the naturally green leaves

⁸ A potential counterexample to this claim was pointed out to us by Nat Hansen. Imagine that Pia has two leaves, one of which is painted green and one of which is naturally green. She cuts both up into many small pieces, mixes the pieces, and then reattaches them using a special kind of glue. According to Hansen, Pia may then hold up one of the two leaves, examine it, and say “*This leaf is greener than that one*” in order to convey the fact that it consists of more naturally green parts than the other leaf. While we agree with the facts, we believe that they actually support our position that this sense of *green* is nongradable. In this example, the comparison is not based on the property named by the adjective: as the example makes clear the concept expressed by *green* remains nongradable. Some other scale must be constructed in order to satisfy the requirements of the comparative morphology; in this case, it is one based on the part structure of the adjective’s nominal arguments. A similar phenomenon can be observed with other nongradable predicates, as in (i):

- (i) The energy used in France is more nuclear than eolic.

We will not provide an analysis of such cases here. But the fact that we need to look to the nouns to build a scale supports the position that the use of *green* involved in this context and Travis’ botanist context, unlike the sense in the artist context, does not itself provide one.

(in a context in which both piles contain objects with the same range of objective color features) shows that, on this use, it is nongradable. The gradable/nongradable distinction is a matter of meaning (cashed out in various ways depending on one's theoretical assumptions, as we will make clear below). It follows, then, that the two utterances of *green* in Travis' story about Pia and her painted leaves involve utterances of distinct terms with distinct meanings, and therefore the sentences in which they are uttered are distinct sentences with distinct conditions for truth. The fact that Pia's utterances of these sentences can be associated with distinct truth values should therefore come as no surprise.

Before we go into the details of the semantic analysis of gradability and our specific proposals about color adjectives, we wish to provide a more complete empirical picture. One of Travis' responses to a proposal like ours is to say that "if words are ambiguous in English, there must be a way of saying just what these ambiguities are; so a fact as to how many ways ambiguous they are" (Travis 1997, p. 90). It is therefore important to be as explicit as we can about what color adjectives mean on their nongradable and gradable uses.

3.2 Classification by color

First, let us consider the nongradable sense. The example of Pia's leaves shows that this meaning should distinguish (among other things) naturally green leaves from painted ones. Let us be more precise about what exactly is relevant about this natural greenness. It is not the case that the relation between the object and the color needs to be natural in the sense of "not man-made." For example, consider a context in which someone is trying to explain the difference in color between the lights on two different traffic signals (one indicating "stop" and the other "go.") It would be extremely odd to do this by pointing to the signal with the red light on and uttering (11a) or (11b) as ways of distinguishing it from the one with the green light on.

- (11) a. #This signal is redder than that.
 b. #That signal is not as red as this one.

Similarly, suppose that our friend Pia is well known as someone who uses only pens that contain blue ink. She finds herself without a pen one day, and asks a colleague if she can borrow one. He hands her a pen from his bag and utters (12a). If Pia tries it out and finds out that in fact her colleague was mistaken and the pen contains black ink, she could very easily hand it back to him with the words in (12b), but an utterance of (12c) would be considered somewhat odd or obnoxious.

- (12) a. Take this one, it's blue.
 b. You're mistaken, that one's black.
 c. ??Thanks, but I want one that's bluer.

What all of these examples share is that having the property denoted by the color adjective is crucially correlated with having some other property or properties which are relevant for some purpose or other. In the case of the green leaves, this might be the property of growing from a plant which will reproduce more leaves of the same color

(vs. one that will not); in the case of the traffic signal, the property of indicating that one must stop or may go; and in the case of the pen, the property of producing blue ink vs. ink of some other color. These correlations constitute the basis for classifying objects: when we classify leaves as green in the botanist's sense, we do so because what really interests us is whether the plant from which they come will grow green leaves or not; arguably, we use the color term because color is an economical and easily observable identifier of this property. The relation between the color property and the correlated property in these cases is not a matter of degree: either the correlation exists or it does not.⁹ For this reason, the color terms as used in these examples are not gradable.

Note that "naturalness" of color is also not a sufficient condition for a color property to be nongradable. The ability of the color property to distinguish objects that have the correlated property from those that lack it is also essential. Consider for instance the use of the color term *blue* to describe one of the naturally occurring colors we attribute to the sky. As far as we know, this use of *blue* is not correlated with the ascription of any single identifiable property to the sky (even though a meteorologist could probably tell us that some such correlation exists—for instance, the lack of a certain amount of cloud cover or humidity, along with the presence of a certain amount of sunlight). Correspondingly, according to our intuitions, all uses of *blue* to describe the sky are gradable (see e.g. (13)) and simply describe its observable color.

- (13) a. The sky is very/completely/perfectly blue.
 b. Today the sky is bluer than it was yesterday.

It is also crucial to point out that there is a certain amount of indeterminacy in how an object that has a classificatory color property (as we will refer to it in the rest of this paper) might manifest physically observable color. For instance, as noted already, a classificatorily green leaf need not be green for its entire existence—it may turn completely brown with age or it may be painted; a red traffic signal need only have a lighted red light; and a blue pen need only have blue ink. This indeterminacy no doubt has its origins in the various relationships that can exist between the manifested color and the correlated property: for example, to signal "stop," what matters is that the traffic signal have one particular light which is red; the color of the rest of the signal is irrelevant. The same holds for the pen. What matters for the leaf is that it is green for a particular period of its life.

If, as we hypothesize, nongradable color terms are used for classifying entities based on their (non)possession of some correlated property, it should come as no surprise that sometimes color terms are chosen to classify entities even when the color is not manifest at all in them, as in the Catalan example in (14).

- (14) Aquest vi és negre.
 this wine is black
 'This wine is red.'

⁹ We will not investigate in detail here whether there are any conditions on the nature of or reason for this correlation (e.g. whether it must be causal; probably it need not be). Rather, we will simply assume that all that is necessary is that the color property be a reliable indicator of the other property or properties which are really of interest.

No Catalan speaker would attribute the color black to *vi negra*; rather, the color term simply serves, along with the terms *blanc* ‘white’ and *rosat* ‘rose-colored’ in Catalan, to create a color-based classification space for wine. Of course it is natural that the color terms chosen should approximate the color of the entities being classified, but [Gärdenfors \(2000, p. 183\)](#), following [Broström \(1994\)](#), suggests that there is a competing tendency to choose the highest possible color terms on [Berlin and Kay \(1969\)](#)’s hierarchy. In such cases, the pressure to choose a more basic color term might exceed the pressure to choose a term that faithfully represents the color of the entities being described.¹⁰ In contrast, as we will show below, gradable color differs sharply from nongradable, classificatory color in allowing radically less variation in the manifestation of the quality and quantity of color that count for the truthful ascription of the color term.

Summarizing, the classificatory reading of color terms consists of a common semantic core with room for a certain amount of indeterminacy. In all cases, the color term denotes the property of manifesting some property which is correlated with the color in question. What exactly that property is and how the color must be manifested are the points on which a certain amount of indeterminacy exists. We will suggest an account of this indeterminacy in [Sect. 4](#).

3.3 Color quantity and color quality

Turning to the gradable interpretation of color adjectives, as happens with many other adjectives, these appear to involve two distinct (sub)readings, which we will refer to as COLOR QUANTITY and COLOR QUALITY (including degree of hue, color saturation, and brightness). The quantity reading expresses a measure of how much of the object is the relevant color; this kind of reading is brought out by proportional modifiers like *completely*, *entirely*, *half*, *part* and so forth, as illustrated in the examples in (15).

- (15) a. Pia painted the leaves part/half/completely green.
 b. Instead of jerseys with blue stripes this year, the team is wearing shirts that are completely/entirely/100% blue.
 c. His pants are half/part green, half/part blue: one leg of each color.

On such a reading, the adjectival property can be associated with a fully closed scale; in other words, in principle there exist both minimal and maximal degrees to which the property can be ascribed ([Rotstein and Winter 2004](#); [Kennedy and McNally 2005](#)). This is the kind of reading that presumably led [Szabó \(2001\)](#) to posit a part structure variable in the semantics of color adjectives, and in fact, proportional modifiers can be used to show that there is reason to believe that Szabó’s part variable (or an analog of it) plays an independent role in the interpretation of color adjectives. Consider the following dialogue, in a context in which speaker A is holding up a red apple.

¹⁰ [Gärdenfors \(2000, p. 182\)](#), again citing [Broström \(1994\)](#) in fact defends the more general view that, as a rule, color terms “aren’t used so much to refer to particular colors as to maintain the color *contrasts* between different referents” (citation from [Broström 1994](#), pp. 101–102). While this view is very much in line with our analysis of nongradable color terms, we think it is too strong for gradable color terms. However, as this point is orthogonal to the main point of this article, we will not discuss it further here.

- (16) A: This apple is completely red.
 B: No, it's not completely red: it's red only on the outside, not on the inside.

A's use of *completely* in (16a) requires *red* to have a quantity interpretation, and it is perfectly imaginable that A means 'completely red on the outside', while B's objection requires an interpretation along the lines of 'completely red throughout'. Given this, we find it plausible to posit that Szabó's part variable is part of the semantic representation of the adjective, or that it could be added via modification of the adjective.

However, this variable does not help us with the color quality reading. This reading involves a measurement of how closely an object's color approximates or diverges from a "center" or prototype. There are different ways to measure approximation to the prototype, which involve not only hue, saturation, brightness, and perhaps other factors based on the physical and perceptual characteristics of color, but also possibly the comparison class which is being used to fix the standard for the color term. The examples in (17), where it is clear that what is being evaluated is the quality of an object's color, rather than the quantity of it that exhibits the color, illustrate this interpretation.

- (17) a. I see that your leaves are all completely painted, but some are greener than others. Try to ensure that the colors are identical.
 b. Your painting is coming along, though it still needs some work: all of the sky is blue, but it isn't blue enough, and the clouds are too white. Try modifying your pigment mixtures.

If the quality reading involves a measurement of proximity to a prototype, we might expect the scales used on these readings to have maximum values corresponding to a perfect match with the prototype. The example in (18a) suggests that this is the case, though the fact that (18b) is not a contradiction shows that the modifier that is used to signal a maximal color quality value is *perfectly* rather than *completely*.¹¹

¹¹ We leave for future research an exploration of exactly how and why different degree modifiers appear to be sensitive to different dimensions (see Sauerland and Stateva 2007 for a recent study of this issue). Note that *perfectly* does not always seem to indicate a maximal value (see for instance Lasersohn 1999 for an analysis of *perfectly* as a slack regulator). For example, (ia) entails that Andrew fits the prototype of height for whatever the relevant standard is in question but does not entail that nobody or nothing else is taller than he is, indicating that the adverb conserves some entailment concerning "perfection" on this use. This perfection entailment no doubt lies behind the oddness of the adverb with adjectives for which some notion of perfection is not obviously relevant, like *open* in (ib).

- (i) a. Don't worry that Andrew isn't right for the basketball player part: He's perfectly tall.
 b. ??The door is perfectly open.

Nonetheless, when the adjective has a scale with a maximal value, even if that value is higher than the standard (as with *clear* used to describe an explanation), modification by *perfectly* does entail a maximal value:

- (ii) a. Your explanation is clear, but it could be clearer.
 b. ??Your explanation is perfectly clear, but it could be clearer.

- (18) a. The baby's eyes were perfectly blue; they couldn't have been bluer.
 b. The baby's shirt was completely blue, but it could have been a bluer shade of blue.

This contrast is paralleled by the contrast in (19) between the nonmaximal degree modifiers *somewhat*, which licenses only the color quality reading, and *half*, which licenses only the color quantity reading.

- (19) a. The car is somewhat green.
 b. The car is half green.

Color terms have generated a lot of attention because of the fact that color quality is not always manifest in exactly the same way, and because the parts of the object that are relevant for deciding the applicability of the color term are not always the same (see e.g. Quine 1960; Lahav 1993). However, once the nongradable reading is teased apart from the gradable reading, color terms begin to look rather more like other adjectives. Consider first color quality. While it is the case that a perfectly red (blushing) face is not the same hue as a perfectly red apple, or that perfectly blue water is probably not the same shade as a perfectly blue sky, these sorts of variations look strikingly similar to the variations we find with any adjective for which a comparison class is relevant, including dimensional adjectives such as *long* or *round*. Obviously the interpretation of such terms includes a context-dependent component, the identity of the comparison class. But it is essential to distinguish the variability of color induced by the variability in comparison class from the variability in color attributable to the nongradable reading. While we have observed above that on the latter reading, it might be possible to find color terms applied to objects which do not in fact manifest the color in question, we would be very surprised to find a case where a gradable color term is used when the color is not manifest at all in the described entities.¹²

Turning to the distribution of color on an object, we again observe that once non-gradable uses of color terms are set aside, judgments are surprisingly consistent as to how much of an object must manifest the color in order for the term to be applicable. Our preliminary observations suggest that in order for a color term to apply to an object on the gradable quantity reading, the color in question must perceptually predominate. To see this, consider the way we attribute color to a piece of clothing. A t-shirt which is entirely white except for a few small red flowers around the neck or three thin pale-grey stripes will pass as white, but a shirt which is white on the entire left-hand side, front and back, and red on the other, will not. The less the color in question predominates on an object, the less likely we are to describe the object as being of that color and the more likely we are to use complex descriptions such as "white with red flowers." If we are right about the consistency of these judgments, the standard for color quantity will be relatively consistent across contexts, contrary to what appears to be the case if nongradable and gradable color attributions are lumped into the same category.

¹² A relevant example in this respect is skin color. Though *white* and *black* and their equivalents in other languages are often used to characterize race, our anecdotal experience indicates that people balk at actually using those terms to characterize the color of an individual's skin.

Finally, note that the well-known puzzle posed by the contrast between the red apple, which is presumably red on the outside, and the red grapefruit, which is presumably red on the inside, looks rather less puzzling if gradable color is distinguished from nongradable color. *Red* is applied to grapefruit in the nongradable sense: it would be bizarre to hold up an uncut, unpeeled red grapefruit and say (20):

(20) This grapefruit is very/completely red.

In contrast, *red* as applied to the exterior of an apple is gradable: sentences such as (21) are perfectly natural.

- (21) a. This apple is very red.
b. This apple is completely red.

The discussion in the literature of *red* as applied to apples implicitly probes intuitions about uncut, unpeeled apples—that is, precisely about the part of the apple that we perceive in such a situation. It is therefore no surprise that what is considered red is just the exterior. But consider a situation in which we present someone with a sliced apple in which the pale-colored flesh is as visible as the red skin or more so. In such a case, while it might not be impossible to use the sentences in (21) to describe the apple, we think it would be unusual to do so. This suggests that as a strong default the part of the object that is relevant for gradable color must be perceptually salient (if amenable to some manipulation in context as suggested in Sect. 3.3).

3.4 Summary

This section has provided evidence that the difference in judgments about the truth of sentences containing color adjectives documented by Travis correlates with distinct interpretations for the adjectives: a nongradable, classificatory meaning on the one hand, and gradable quality/quantity meanings on the other. Although the latter two meanings can be clearly differentiated, we will argue in the next section that the difference between them is not one of semantic type: in both cases, the adjective denotes a function from individuals to degrees. In contrast, when an adjective takes on a classificatory meaning it denotes a nongradable property of individuals. The semantic difference between gradable and nongradable color is therefore both a matter of meaning (content) and semantic type, and so constitutes a case of ambiguity. If this is correct, then Travis' argument against a compositional theory of truth conditions based on sentences with color adjectives disappears.

4 A semantics for color terms

In the previous section, we argued that color adjectives are ambiguous between a nongradable, classificatory interpretation and gradable quantity/quality interpretations. Intuitively, however, these meanings have something in common, and we want our analysis to capture that fact. To do this we will appeal to the fact that color terms are also nouns and we will provide an analysis of adjectival denotations in terms of (what we take to be) the more basic nominal ones.

Color nouns are routinely used to refer to colors, and as exemplified in (22), they are mass nouns, as their ability to appear in bare singular form and their compatibility with mass quantifiers shows.

- (22) a. Green was a surprising choice for the color of the dining room.
b. There was a little/not much green in the carpet.

We will assume following Carlson (1977) and others that, like other mass nouns, they can denote the kinds of entities that they describe, as in (23), where **green** is a constant of type e .

$$(23) T(\text{green}_N) = \mathbf{green}$$

The instances of the kinds color terms describe—and indeed possibly the definition of the kinds themselves—are vague because colors are continuous, but in this respect the situation is no different from that of many other nouns, such as *village*, *town* or *city*. However, this sort of vagueness, which is independent of the gradable/nongradable distinction as well as the distinction between color quantity and color quality, will not concern us here. When the noun is changed into an adjective, the semantics of the noun can feed into the semantics of the resulting adjective, but it will do so in different ways for the gradable and nongradable readings. Let us consider the latter first.

As argued above, a classificatory color adjective applies truthfully to its argument just in case that argument has some other property which is correlated with the color named by the nominal form. We represent this meaning as shown in (24), where the correlation relation is indicated as **cor** and the correlated property is the value of a free variable P_i .

$$(24) T(\text{green}_A^{\text{nongr}}) = \lambda x. P_i(x) \wedge \mathbf{cor}(P_i, \mathbf{green})$$

This denotation is type $\langle e, t \rangle$, and specifies no generalized satisfaction conditions for nongradable *green* beyond the requirement that an object have the classifying property and that the property be correlated with the color **green**. We assume that the correlation relation generally requires the object to also manifest the relevant color somehow, somewhere, at some point in time, but this variability, together with the underspecification of the color-correlated property, introduces a certain indeterminacy into the semantics for nongradable color that we have seen to be consistent with the facts (see the discussion in Sect. 3.2).¹³ In this sense, our proposal bears some affinity with the fully indexical analysis in Rothschild and Segal (2009); the difference is that in our account, the “classification by color” meaning that (24) is designed to represent is an aspect of the truth conditions that remains constant no matter how we fix the value of P_i , the correlated property.

For the gradable interpretations, we treat color adjectives like other gradable predicates. Specifically, we follow Kennedy (1999) and analyze them as denoting functions from objects to degrees, in this case, degrees that represent the extent to which the

¹³ Although we treat P_i as a free variable here, we assume that its value is conventionally determined by features of the object denoted by the color adjective’s argument. For example, *green* as applied nongradably to traffic lights will be associated with a different value for P_i than will the same adjective applied nongradably to leaves. This sort of sensitivity is common in the semantics of modification; a promising way of treating it is presented in Asher (2007).

object manifests the color named by the noun. Based on the empirical observations in Sect. 3.3, we assume two variants of this measure function, one that measures how much of the object manifests the color (the quantity reading) and one that measures how closely the object's manifestation of the color approximates the appropriate prototype (the quality reading). We represent these denotations in (25a–b), respectively.¹⁴

- (25) a. $T(\text{green}_A^{\text{quant}}) = \lambda x. \mathbf{quant}(\mathbf{green})(x)$
 b. $T(\text{green}_A^{\text{qual}}) = \lambda x. \mathbf{qual}(\mathbf{green})(x)$

The gradable denotations of color adjectives are type $\langle e, d \rangle$, and just like other gradable adjectives on Kennedy's analysis, they are converted into properties of individuals by combining with degree morphology. Different degree morphemes introduce different kinds of STANDARDS that determine whether the property in question is held in sufficient degree for the predicate to truthfully apply to its argument. For example, the (unmodified) positive form involves a null degree morpheme *pos* that introduces a relation to a contextual standard of comparison, while the degree term *very* introduces a higher standard, one that is determined by a comparison class consisting of the individuals that satisfy the (unmarked) positive form. Finally, comparative morphology introduces a standard that corresponds to the degree to which the adjectival property is held by another individual, typically introduced by a comparison clause. These various denotations are provided in (26) (cf. Kennedy 1999, 2007; Kennedy and McNally 2005).¹⁵

- (26) a. $T(\text{pos}) = \lambda g_{\langle e, d \rangle} \lambda x. g(x) > \mathbf{stnd}(g)$
 b. $T(\text{very}) = \lambda g_{\langle e, d \rangle} \lambda x. g(x) > \mathbf{stnd}(g_{\{x | \text{pos}(g)(x)\}})$
 c. $T(\text{--er . . . than}) = \lambda g_{\langle e, d \rangle} \lambda x \lambda y. g(y) > g(x)$

Returning to Travis' original example, repeated again in (27a), we can now associate it with three distinct propositions, represented as in (27b–d).

- (27) a. The leaves are green.
 b. $P_i(\mathbf{the\ leaves}) \wedge \mathbf{cor}(P_i, \mathbf{green})$
 c. $\mathbf{quant}(\mathbf{green})(\mathbf{the\ leaves}) > \mathbf{stnd}(\mathbf{quant}(\mathbf{green}))$
 d. $\mathbf{qual}(\mathbf{green})(\mathbf{the\ leaves}) > \mathbf{stnd}(\mathbf{qual}(\mathbf{green}))$

(27b) is true just in case the property that P_i ends up being valued to is true of the leaves and this property is (conventionally) correlated with the color green; (27c) is true just in case the quantity of the leaves' greenness exceeds a standard of comparison for this measure (in the context of utterance); and (27d) is true just in case the quality of the leaves' greenness exceeds a standard of comparison for this kind of measure (in the context). It is therefore no surprise that different utterances of the string of words in (27a) are "compatible with different ways of being true".

¹⁴ Though we have represented them separately here, a general theory of gradable adjective meaning could treat quality and quantity as two of a limited set of values for a DIMENSION parameter that would be specifiable in different ways according to the gradable property in question. Viewed this way, this aspect of the adjective's meaning is context dependent, and not due to a genuine lexical ambiguity.

¹⁵ Note that the standard introduced by *pos* (and, consequently, *very*) is yet another source of vagueness in the semantics of gradable color terms, distinct from the others we have seen so far. See Kennedy (2007) and the references cited therein for extensive discussion of the vagueness associated with standards of comparison.

5 Conclusions

Travis' argument was based on the assumption that there was no ambiguity of any kind in examples like (1), repeated here:

(28) The leaves are green. (= (1))

However, we have shown that this assumption is incorrect by demonstrating a correlation between the true-false judgments in Travis contexts and the presence/absence of gradability. This distinction in the semantics of color adjectives in the two uses supports the hypothesis that we have distinct terms in each of the two relevant utterances of (1), with their own distributional and interpretive features. In other words, the different utterances of (1) correspond to utterances of distinct sentences, which convey distinct propositions, and there is no expectation that they should have the same truth conditions.

It should also now be evident why an analysis like Szabó's (2001) is insufficient to provide a full explanation of the linguistic facts, even though we showed via (16) that there may be independent evidence to posit a part variable in the semantic representation of color adjectives (at least optionally). The problem is that a part variable cannot capture the distinction that we have seen between gradable and nongradable senses of color terms. We might try to derive the nongradable interpretation by maximizing over the part variable—requiring it to pick out all parts of the argument of the adjective. But this would equate the nongradable sense of e.g. *green* with *completely green*, which is not correct. Such an analysis would not explain the Travis facts (see the discussion at the end of Sect. 2.1), nor would it explain the data in (10), because the fact that an unmodified gradable adjective picks out a maximal value in a particular context does not rule out the possibility of degree morphology in the same context. This is illustrated by the examples in (29), which use the “maximum standard” adjective *flat* (Kennedy and McNally 2005; Kennedy 2007). Even if we are in a context in which we have agreed that (29a) is true of one tabletop and false of a second one, an utterance of any of the various modified forms in (29b–e) would also be felicitous.¹⁶

- (29) a. This tabletop is flat.
 b. This tabletop is perfectly flat.
 c. It's even flatter than that one.
 d. That tabletop is not as flat as this one.
 e. That tabletop is not very flat.

¹⁶ A similar objection applies to the analysis of color adjectives in Wheeler (1972), who notes that they can have ‘absolute’ interpretations, and attempts to account for such meanings by assuming that they have comparison class variables that are obligatorily fixed to the universe of discourse. On this view (what we have been calling) nongradable *green* means something like *green compared to anything*.

However, like Szabó's analysis, this proposal fails to explain the data in (10). In Wheeler's account, the semantics of degree constructions involves manipulation of a comparison class (e.g., *a is more F than b* is equivalent to *a is F relative to a comparison class consisting of b*), and even adjectives whose comparison classes are conventionally fixed to the domain of discourse when they appear without modifiers allow for this. This is illustrated by the data in (29), since the absolute adjective *flat*, in Wheeler's account, receives exactly the same semantic analysis as a color adjective (see Wheeler 1972, pp. 330–333).

Szabó's analysis also has nothing to say about the differences between the color quality and color quantity readings, suggesting that a more general analysis of color terms is called for, one which is embedded within an independently motivated theory of the semantics of adjectives and from which the Travis facts should follow directly.

We close with some final comments on the indexical response to Travis articulated by [Rothschild and Segal \(2009\)](#). The three types of color adjective denotations we have argued for could of course be derived on such an account: if color adjectives are indexical predicates, then the denotations we have posited are presumably among their possible valuations. But there are two disadvantages to such an account. First and most importantly, we have no explanation for why we systematically find these particular denotations, and only these denotations. As we argued in Sect. 2.2, appeals to psychology are not sufficient to explain this fact, but this is the only recourse open to such an approach.

In contrast, on our analysis both the gradable and nongradable denotations can be seen as specific instantiations of more general grammatical regularities. On their gradable denotations, color terms behave much like many other gradable adjectives, which can be interpreted either with respect to scales that are determined by the part structure of the object to which they are ascribed (yielding a quantity reading) or with respect to a qualitative dimension. For example, the adjective *wet* has both quality and quantity readings, as illustrated in (30a) and (30b), respectively.

- (30) a. The tattoo must be completely wet (though not necessarily soaking) in order for it to transfer properly to the child's hand.
 b. After the rain the playground was very wet, though the children were able to play in a dry area under a tree.

Similarly, the gradable/nongradable ambiguity in color terms is paralleled by similar, if not identical, ambiguities, in other classes of adjectives such as nationality adjectives (e.g. *French*, *Siberian*), other so-called relational adjectives (e.g. *presidential*, *musical*), and other denominal adjectives (e.g. *wooden*, *woolen*). Our approach to color adjectives brings them in line with these other expressions, and allows for the possibility of eventually capturing these correlations in terms of a general theory of the lexical and compositional semantics of adjectives, avoiding recourse to the power of indexicality. This theory, which will most likely be grounded in facts about human psychology or the way the world is, will also almost certainly reflect facts that are particular to the way natural language works.

The second disadvantage of a strictly indexical account is that, while it rightly recognizes the context dependence in the interpretation of color terms, it fails to distinguish the particular kinds of context dependence that they manifest. Once context dependency is identified in meaning, it's a fast and slippery slope to the conclusion that everything is indexical. However, easy-to-overlook linguistic data from expressions like degree modifiers show that humans do limit the semantic space of possibilities and grammaticize those limits to a certain degree. Such observations underscore the relevance of the results of linguistic research for debates such as those raised by the Travis facts.

Acknowledgements Thanks to Gemma Boleda, Nat Hansen, Josep Macià, and Daniel Rothschild for helpful discussion, and to participants in the Lille International Conference on Adjectives, the Universitat de Barcelona LOGOS Semantics Reading Group, and the University of Chicago Workshop on Semantics and Philosophy of Language for feedback. This paper is based upon work supported by the National Science Foundation under Grants No. 0094263, 0618917 and 0620247, by the Spanish Ministerio de Ciencia e Innovación under Grant No. HUM2007-60599/FILO, and by a grant from the Catalan Department of Universities, Research, and Information Society.

References

- Asher, N. (2007). A web of words: Lexical meaning in context. Ms., CNRS-IRIT and U. Texas Austin.
- Berlin, B., & Kay, P. (1969). *Basic color terms: Their universality and evolution*. Berkeley, CA: University of California Press.
- Broström, S. (1994). *The role of metaphor in cognitive semantics*. Lund: Lund University Cognitive Studies 31.
- Carlson, G. N. (1977). *Reference to kinds in English*. PhD thesis, University of Massachusetts, Amherst.
- Chierchia, G. (1995). Individual-level predicates as inherent generics. In G. Carlson & F. J. Pelletier (Eds.), *The generic book* (pp. 176–223). Chicago, IL: University of Chicago Press.
- Gärdenfors, P. (2000). *Conceptual spaces: On the geometry of thought*. Cambridge, MA: MIT Press.
- Grinder, J., & Postal, P. (1971). Missing antecedents. *Linguistic Inquiry*, 2, 269–312.
- Hankamer, J., & Sag, I. (1976). Deep and surface anaphora. *Linguistic Inquiry*, 7, 391–428.
- Kennedy, C. (1999). *Projecting the adjective: The syntax and semantics of gradability and comparison*. New York: Garland (1997 UCSC PhD thesis).
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable predicates. *Linguistics and Philosophy*, 30(1), 1–45.
- Kennedy, C., & McNally, L. (2005). Scale structure and the semantic typology of gradable predicates. *Language*, 81(2), 345–381.
- Kratzer, A. (1995). Stage-level and individual-level predicates. In G. Carlson & F. Pelletier (Eds.), *The generic book* (pp. 125–175). Chicago, IL: University of Chicago Press.
- Lahav, R. (1993). The combinatorial-connectionist debate and the pragmatics of adjectives. *Pragmatics and Cognition*, 1, 71–88.
- Lakoff, G. (1970). A note on vagueness and ambiguity. *Linguistic Inquiry*, 1(3), 357–359.
- Laserson, P. (1999). Pragmatic halos. *Language*, 75(3), 522–551.
- Predelli, S. (2005). Painted leaves, context, and semantic analysis. *Linguistics and Philosophy*, 28, 351–374.
- Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.
- Rothschild, D., & Segal, G. (2009). Indexical predicates. *Mind and Language*, 24(4), 467–493.
- Rotstein, C., & Winter, Y. (2004). Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics*, 12, 259–288.
- Sainsbury, R. M. (2001). Two ways to smoke a cigarette. *Ratio*, 14(4), 386–406.
- Sauerland, U., & Stateva, P. (2007). Scalar vs. epistemic vagueness: Evidence from approximators. In M. Gibson & T. Friedman (Eds.), *Proceedings of SALT 17* (pp. n–m). Ithaca, NY: CLC Publications.
- Szabó, Z. (2001). Adjectives in context. In I. Kenesei & R. M. Harnish (Eds.), *Perspectives on semantics, pragmatics and discourse: A festschrift for Ferenc Kiefer* (pp. 119–146). Amsterdam: John Benjamins.
- Travis, C. (1985). On what is strictly speaking true. *Canadian Journal of Philosophy*, 15(2), 187–229.
- Travis, C. (1994). On constraints of generality. *Proceedings of the Aristotelian Society*, 94(2), 165–188.
- Travis, C. (1997). Pragmatics. In B. Hale & C. Wright (Eds.), *A companion to the philosophy of language* (pp. 87–106). Oxford: Blackwell.
- Wheeler, S. (1972). Attributives and their modifiers. *Noûs*, 6(4), 310–334.
- Zwicky, A., & Sadock, J. (1975). Ambiguity tests and how to fail them. In J. P. Kimball (Ed.), *Syntax and semantics* (Vol. 4, pp. 1–36). New York: Academic Press.