

GRADABLE ADJECTIVES DENOTE MEASURE FUNCTIONS, NOT PARTIAL FUNCTIONS*

Christopher Kennedy
Northwestern University
kennedy@nwu.edu

This paper uses the distribution and interpretation of antonymous adjectives in comparative constructions to argue that an empirically and explanatorily adequate semantics for gradable adjectives must introduce abstract representations, or ‘degrees’, into the ontology. I begin with an overview of the basic assumptions of analyses that do not make reference to scales and degrees. I then turn to a discussion of the empirical data, and I demonstrate that such approaches do not support a principled explanation of the facts, but a degree-based account does.

1. Gradable Adjectives and the Problem of Vagueness

A fundamental problem for the semantic analysis of gradable adjectives like *tall*, *long*, and *expensive* is that many of the sentences in which they occur are vague. (1), for example, could be true in one context and false in another.

(1) The Mars Pathfinder mission was expensive.

In a context in which the discussion includes all objects that have some cost value associated with them, (1) would most likely be judged true, since the cost of sending a spacecraft to Mars is far greater than the cost of most things (e.g., nails, dog food, a used Volvo, etc.). In a context in which only missions involving interplanetary exploration are salient, however, then (1) would probably be judged false, since a unique characteristic of the Mars Pathfinder mission was its low cost compared to other projects involving the exploration of outer space (see Sapir 1944, McConnell-Ginet 1973, Kamp 1975, Klein 1980, Ludlow 1989, Kennedy 1999 and others for relevant discussion).

What this example shows is that the criteria for deciding whether an utterance of a sentence of the form ‘*x is φ* ’ is true (where φ is a gradable adjective) may vary according to factors external to the adjective, such as the meaning of *x* and features of the context of utterance. Determining the truth of ‘*x is φ* ’ in a particular context requires figuring out what these criteria are, and then making a judgment of whether *x* ‘counts as’ φ in that context. A basic requirement of a theory of the semantics of gradable adjectives, then, is to both provide a means of making this judgment, ensuring that sentences like (1) have

definite interpretations, and to allow for variability of interpretation across contexts.

One approach to this problem, first formalized in Seuren (1973) and Cresswell (1976) but since adopted in some form by many analyses of the semantics of gradable adjectives, meets this requirement by constructing an abstract representation of measurement and defining the interpretation of gradable adjectives in terms of this representation (see e.g., Hellan 1981, Hoeksema 1983, von Stechow 1984a,b, Heim 1985, Bierwisch 1989, Pinkal 1989, Moltmann 1992a, Gawron 1995, Rullmann 1995, Hendriks 1995, and Kennedy 1999; see Klein 1991 for general discussion of such approaches). This abstract representation, or *scale*, can be construed as a set of objects totally ordered along some dimension (such as *height*, *width*, *density* and so forth; the dimension value has the effect of distinguishing one scale from another), where each object represents a measure, or *degree*, of ' φ -ness'. The introduction of scales and degrees into the ontology makes it possible to analyze predications involving gradable adjectives as relations between objects in their domains and degrees on a scale, which in turn provides the basis for an account of vagueness. A sentence of the form ' x is φ ' is taken to mean that the degree to which x is φ is at least as great as some other degree d_s on the scale associated with φ that identifies a *standard* for φ . The semantic function of the standard-denoting degree is to provide a means of separating those objects for which the statement ' x is φ ' is true from those objects for which ' x is φ ' is false (see Bartsch and Vennemann 1973, Cresswell 1976, Siegel 1979, Klein 1980, von Stechow 1984a, Ludlow 1989, Bierwisch 1989, and Kennedy 1999). The problem of vagueness is thus recast as the problem of determining the value of the standard in a particular context.

For example, a sentence like (1), on this view, is true just in case the degree that represents the cost of the Mars Pathfinder mission is at least as great as the standard value for 'expensive' in the context of utterance.¹ In a context in which all objects in the domain of expensive are relevant, the standard value would fall below the degree that represents the cost of the Mars Pathfinder mission, and (1) would be false. In a context in which only interplanetary expeditions are considered, however, the standard value would exceed the cost of the Pathfinder mission, and (1) would work out to be true.

A question that this type of approach must deal with, however, is whether the introduction of scales and degrees into the ontology is really necessary. Is it possible instead that the semantic properties of gradable adjectives can be equally well explained in terms of some alternative mechanisms that do not make reference to abstract objects? This question is of particular importance, because such alternative analyses have been developed (see for example McConnell-Ginet 1973, Kamp 1975, Klein 1980, 1982, Van Benthem 1983, Larson 1988, and Sánchez-Valencia 1994). These approaches build vagueness directly into the meaning of a gradable adjective, by characterizing its meaning in terms of a family of functions from individuals to truth values, some of which may be partial. In any context of use, a particular function from this family must be selected in order to achieve a definite interpretation, but no reference to degrees or other abstract objects is required.

The goal of this paper is to address this question by investigating a set of facts that provide unusual insight into the empirical consequences of different assumptions about the nature and ontological status of scales and degrees.

Specifically, I will argue that the distribution and interpretation of antonymous pairs of adjectives in comparatives indicate that an explanatorily adequate semantics of gradable adjectives must make reference to abstract representations of measurement, since it is precisely the properties of such representations that are crucial for explaining the observed facts in this domain. Before going into the data, however, I will present an overview of the basic assumptions of theories that do not make reference to abstract representations of measurement.

2. Capturing Vagueness without Degrees

2.1 Gradable Adjectives

For perspicuity, I focus in this section on the analysis of gradable adjectives articulated in Klein 1980, 1982, but my remarks hold equally of the other analyses mentioned above. The starting point of this approach, which I will refer to as the ‘partial function analysis’, is the assumption that gradable adjectives are of the same semantic type as non-gradable adjectives: they denote functions from objects to truth values. Gradable adjectives are distinguished from non-gradable adjectives (and other predicative expressions) in that their domains are partially ordered according to some property that permits grading, such as *cost*, *temperature*, *height*, or *brightness* (this property corresponds to the concept of ‘dimension’ that scalar approaches use to distinguish one scale from another). Klein (1980, 1982), building on Kamp 1975, makes a second distinction between gradable and non-gradable adjectives: the latter always denote complete functions from individuals to truth values, but the former can denote partial functions from individuals to truth values. In other words, non-gradable adjectives like ‘hexagonal’ and ‘Croatian’ always denote functions that return a value in $\{0,1\}$ when applied to objects in their domains, but gradable adjectives like ‘dense’, ‘bright’, and ‘shallow’ can denote functions that return 0, 1 or no value at all when applied to objects in their domains.

The interpretation of a proposition with a gradable adjective as the main predication can be stated as follows. First, assume as above that the domain of a gradable adjective is partially ordered according to some dimension. A gradable adjective φ in a context c can then be analyzed as a function that induces a tripartite partitioning of its (ordered) domain into: (i) a positive extension ($pos_c(\varphi)$), which contains objects above some point in the ordering (objects that are definitely φ in c), (ii) a negative extension ($neg_c(\varphi)$), which contains objects below some point the ordering (objects that are definitely not φ in c), and (iii) an ‘extension gap’ ($gap_c(\varphi)$), which contains objects that fall within an indeterminate middle, i.e., objects for which it is unclear whether they are or are not φ in c (cf. Sapir’s (1944) ‘zone of indifference’). The net effect of these assumptions is that the truth conditions of a sentence of the form ‘ x is φ ’ in a context c can be defined as in (2).

- (2) a. $\|\varphi(x)\|^c = 1$ iff x is in the positive extension of φ at c ,
 b. $\|\varphi(x)\|^c = 0$ iff x is in the negative extension of φ at c , and
 c. $\|\varphi(x)\|^c$ is undefined otherwise.

The partitioning of the domain into positive and negative extension and extension gap is context-dependent, determined by the choice of comparison class. Roughly speaking, a comparison class is a subset of the domain of discourse that is determined to be somehow relevant in the context of utterance, and it is this subset that is supplied as the domain of the function denoted by the adjective.² The role of the comparison class can be illustrated by considering an example like (3).

(3) Erik is tall.

If the entire domain of discourse were taken into consideration when evaluating the truth of (3), then it would turn out to be either false or undefined, since relative to mountains, redwoods, and skyscrapers, humans fall at the lower end of an ordering along a dimension of height. As a result, the individual denoted by *Erik* would be at the lower end of the ordered domain of the adjective, and so would fall within the negative extension of tall (or possibly in the extension gap). When attention is restricted to humans, however, then a comparison class consisting only of humans is used as the basis for the partitioning of the domain of tall, and the truth or falsity of depends only on the position of Erik in this smaller set.

An important constraint on the construction of a comparison class is that it must preserve the original ordering on the domain, in order to avoid undesirable entailments. For example, consider a context in which the ordering on the domain of tall is as in (4).

(4) $D_{tall} = \langle \dots, Nadine, Bill, Aisha, Chris, Tim, Frances, Polly, Erik, \dots \rangle$

If no restrictions were placed on the construction of a comparison class from D_{tall} , then the ordered set K_{tall} in (5) would be a possible comparison class, allowing for a partitioning of the domain as shown in (6).

(5) $K_{tall} = \langle \dots, Aisha, Frances, Polly, Nadine, \dots \rangle$

- (6) a. $pos_c(tall) = \langle Polly, Nadine \rangle$
 b. $neg_c(tall) = \langle Aisha \rangle$
 c. $gap_c(tall) = \langle Frances \rangle$

In this context, (7) would be true while (8) would be false, a result which is inconsistent with our intuitions if the actual ordering on the domain of tall is as in (4).

(7) Nadine is tall.

(8) Frances is tall.

This undesirable result is avoided by invoking the *Consistency Postulate*, informally defined in (9) (see Klein 1980, 1982, van Benthem 1983, Sanchez-Valencia 1994, and below), which requires any partitioning of a subset of the domain of a gradable adjective to preserve the original ordering on the entire domain.

(9) *Consistency Postulate* (informal)

For any context in which ‘*a is φ* ’ is true and $b \geq_{\varphi} a$, then ‘*b is φ* ’ is also true, and for any context in which ‘*a is φ* ’ is false, and $a \geq_{\varphi} b$, then ‘*b is φ* ’ is also false.³

A consequence of this condition is that (6) is not a possible partitioning. Since the partitioning indicated in (6) makes (7) true, and *Frances* \geq_{tall} *Nadine* with respect to the original ordering in (4), the Consistency Postulate requires it to also be the case that (8) is true.

This discussion brings into focus the fundamental ideas underlying the partial function analysis. Given any set of objects partially ordered along a dimension δ , it is possible to define a family of (possibly partial) functions that induce a partitioning on the set in accord with the Consistency Postulate. In effect, the partial function analysis claims that the interpretation of a gradable adjective with dimensional parameter δ is a value selected from this family of functions, which may vary from context to context. The vagueness of gradable adjectives stems from the fact that in any context of use, it is necessary to choose some function from this family as the interpretation of the adjective in that context.

2.2 Comparatives

The intuitions underlying the Consistency Postulate provide the foundation for an analysis of comparatives (see McConnell-Ginet 1973, Kamp 1975, Klein 1980, 1982, van Benthem 1983, Larson 1988a, Sánchez-Valencia 1994). Consider, for example, (10).⁴ In addition, see Larson 1988 for some refinements of the basic analysis developed to handle the interpretation of quantificational expressions in the comparative clause (the complement of *than* or *as*).

(10) Jupiter is larger than Saturn (is).

Given the conditions imposed by the Consistency Postulate, it follows that if there is a context that makes the proposition expressed by ‘*Jupiter is large*’ true but makes ‘*Saturn is large*’ false, then it must be the case that the object denoted by *Jupiter* is ordered above the object denoted by *Saturn* with respect to the ordering on the domain of *large*, i.e., it must be the case that Jupiter is larger than Saturn.

This analysis can be made precise by building on the observation made at the end of the previous section that the interpretation of a gradable adjective in a context *c* is a member of a family of functions that partition a partially ordered set in accord with the Consistency Postulate. Specifically, we can introduce a set of *degree functions* that apply to a gradable adjective and return some member of this family; in particular, following Klein 1980, we can assume that the result of applying a degree function to a gradable adjective is always a complete function. (In Klein’s analysis, the denotations of *very*, *fairly*, and other degree modifiers are taken from the set of degree functions.) The underlying idea is that a degree function performs the role normally played by context: it fixes the denotation of the adjective, ultimately determining how the domain is to be partitioned. The difference is that all of the partitionings induced by a degree function are bipartite: none contain an extension gap.

Once we have degree functions, the Consistency Postulate can be restated more formally as in (11), where G is the set of gradable adjective meanings, D is the domain of discourse, and Deg is the set of degree functions (cf. Klein 1982:126).

(11) *Consistency Postulate*

$\forall \varphi \in G, a, b \in D, c \in C, d \in Deg:$

$[\| (d(\varphi))(a) \| ^c = 1 \ \& \ b \geq_{\varphi} a] \rightarrow [\| (d(\varphi))(b) \| ^c = 1]$ and

$[\| (d(\varphi))(a) \| ^c = 0 \ \& \ a \geq_{\varphi} b] \rightarrow [\| (d(\varphi))(b) \| ^c = 0]$

The effect of the Consistency Postulate is to ensure that the only admissible degree functions are those that induce partitionings of the domain of a gradable adjective in a way that is consistent with the inferences discussed in the previous section. Given this constraint, the interpretation of comparatives can be straightforwardly formalized in terms of quantification over degree functions. A typical comparative of the form *a is more φ than b* is assigned the truth conditions in (12) (the formalism adopted here is most similar to that in Klein 1982).

(12) $\exists d[\| (d(\varphi))(a) \| \ \& \ \sim \|(d(\varphi))(b)\|]$

For illustration, consider the analysis of (10), which has the logical representation in (13).

(13) $\exists d[\| (d(large))(Jupiter) \| \ \& \ \sim \|(d(large))(Saturn)\|]$

According to (13), (10) is true just in case there is a function that, when applied to *large*, induces a partitioning of the predicate's domain so that the positive extension includes *Jupiter*, while the negative extension contains Saturn. Assuming the domain of *large* to be as in (14) (limiting the domain to the planets in the solar system), (10) is true, because there is a partitioning of the domain of *large* such that Jupiter is in the positive extension and Saturn is in the negative extension, namely the one shown in (15). (To distinguish the partitioning introduced by a degree function from the context-dependent partitioning associated with the absolute construction, I have represented the positive and negative extensions induced by a particular degree function d as $pos_d(\varphi)$ and $neg_d(\varphi)$, respectively.)

(14) $D_{large} = \langle Pluto, Mercury, Mars, Venus, Earth, Neptune, Uranus, Saturn, Jupiter \rangle$

(15) a. $pos_d(large) = \langle Jupiter \rangle$

b. $neg_d(large) = \langle Pluto, Mercury, Mars, Venus, Earth, Neptune, Uranus, Saturn \rangle$

Since the possible values of the function d must satisfy Consistency Postulate, partitionings such as (16) are impossible, and we derive the desired result that (10) entails that for any context, *Jupiter* is ordered above *Saturn* in the domain of *large*; i.e., that Jupiter is larger than Saturn is.

- (16) a. $pos_d(\textit{large}) = \langle \textit{Uranus}, \textit{Saturn} \rangle$
b. $neg_d(\textit{large}) = \langle \textit{Pluto}, \textit{Mercury}, \textit{Mars}, \textit{Venus/Earth}, \textit{Neptune}, \textit{Jupiter} \rangle$

3. Comparison and Polar Opposition

The previous section demonstrated that it is possible to construct a descriptively adequate semantics for gradable adjectives and comparatives without making reference to scales and degrees. The purpose of this section is to introduce a set of data involving the distribution and interpretation of antonymous adjectives in comparatives that indicate that this analysis cannot be maintained. Instead, these facts provide strong evidence that an explanatorily adequate semantic analysis of gradable adjectives must in fact make reference to abstract representations of measurement—scales and degrees.

3.1 Cross-polar Anomaly

The first set of empirical facts that will form the basis for my argument is a phenomenon that I refer to as cross-polar anomaly (CPA), which is exemplified by the sentences in (17) (see Hale 1970, Bierwisch 1989, and Kennedy 1997).

- (17) a. ? Alice is shorter than Carmen is tall.
b. ? *The Brothers Karamazov* is longer than *The Idiot* is short.
c. ? The Mars Pathfinder mission was cheaper than the Viking mission was expensive.
d. ? New York is dirtier than Chicago is clean.
e. ? A Volvo is safer than a Fiat is dangerous.

These sentences demonstrate that comparatives formed out of so-called ‘positive’ and ‘negative’ pairs of adjectives are semantically anomalous.⁵ This anomaly cannot be accounted for in terms of syntactic ill-formedness: structurally identical comparatives in which both adjectives have the same polarity, such as those in (18), are perfectly well-formed.

- (18) a. The space telescope is longer than it is wide.
b. After she swallowed the drink, Alice discovered that she was shorter than the doorway was low.

Given the acceptability of examples like these, we can conclude that the factors underlying cross-polar anomaly involve the interaction of the semantics of positive and negative adjectives and the semantics of the comparative construction (see Kennedy 1999b for extensive argumentation in support of this conclusion).

The second set of facts that I will focus on contains examples that are superficially similar to examples of CPA, but are not anomalous.⁶ These examples, which I will refer to as comparison of deviation (COD) constructions, are illustrated by the naturally occurring sentences in (19).

- (19) a. [The Red Sox] will be scrutinized as closely as the Orioles to see

- whether they are any more legitimate than the Orioles are fraudulent.
(*New York Times*, Summer 1998)
- b. Grace especially had a forgettable playoff series that won't soon be forgotten. Grace was as cold as he was hot in the 1989 playoffs.
(*Chicago Tribune*, October 4, 1998, Section 3, p. 6)
 - c. I can still remember the sound it made, a lovely special sound, as light and thin as the clothes were solid and heavy.
(Dibdin, M.: 1989, *Ratking*, Bantam Crime, New York, p. 178)

COD constructions have two characteristics that distinguish them from standard comparatives. First, examples of COD compare the relative extents to which two objects deviate from some standard value associated with the adjective (cf. Bierwisch 1989:220). The meaning of the comparative in (19a), for example, can be paraphrased as in (20).

(20) The degree to which the Red Sox exceed a standard of legitimacy is greater than the degree to which the Orioles exceed a standard of fraudulence.

In contrast, standard comparatives and equatives compare the absolute projections of two objects on a scale. The most natural paraphrase of the equative construction in (21), for example, is (22).

(21) It was a squarish hole, as deep as a ten-story building is tall, cut down into the hard and uncooperative earth.

(Reynolds, W.J., 'The Lost Boys', in Hillerman, T.: 1994, *The Mysterious West*, Harper Collins, New York, p. 223)

(22) The depth of the hole is at least as great as the height of a ten-story building.

Second, unlike typical comparatives, COD constructions entail that the properties predicated of the compared objects are true in the absolute sense. (23a), for example, is contradictory, but (23b) is not.

- (23) a. The Red Sox are more legitimate than the Orioles are fraudulent, but they're not legitimate.
- b. The hole is deeper than a two-year old is tall, but it's not deep.

This property is clearly related to the interpretation of COD, and follows straightforwardly in a scalar analysis (though as we will see below, it is problematic for the partial function analysis). Since the truth of an expression of the form ' x is φ ' in such an analysis is determined by checking whether the degree to which x is φ exceeds an appropriate standard value (see the discussion of this point in section 1), the fact that comparison of deviation constructions compare the degrees to which two objects exceed their respective standard values derives the observed entailment patterns.

An important point to make about comparison of deviation is that in comparatives of the sort under discussion here—examples constructed out of positive and negative pairs of adjectives—the COD interpretation is the only interpretation available. This is most clearly illustrated by equative constructions such as (19b)-(19c) and (24).

(24) The Cubs are as old as the White Sox are young.

This sentence cannot mean that the (average) age of the players on the Cubs is the same as the (average) age of the players on the White Sox, which is what a standard equative interpretation would give us (cf. (21)-(22) above). It can only mean that the degree to which the average age of the Cubs exceeds a standard of oldness (for baseball teams) is the same as the degree to which the average age of the White Sox exceeds a standard of youngness (for baseball teams).

This fact is extremely important because it highlights the fact that comparatives constructed out of antonymous pairs of adjectives are anomalous on what we can call the ‘standard’ comparative interpretation—one in which the absolute degrees to which two objects possess some gradable property is being compared. It follows that a minimal requirement of descriptive adequacy for any semantic analysis of gradable adjectives and comparatives is that it must entail that comparatives formed out of antonymous pairs should be acceptable *only* on a COD interpretation. As we will see in the next section, this is exactly where the partial function analysis fails.

3.2 The Problem of Cross-polar Anomaly for a Klein-style Analysis

Although Klein (1980) does not explicitly discuss the differences between antonymous pairs of positive and negative adjectives such as *tall~short*, *clever~stupid*, and *safe~dangerous*, a natural approach to adjectival polarity within a partial function analysis is to assume, building on the observations about the logical properties of gradable adjectives discussed in section 2.1, that the domains of antonymous pairs are distinguished by their orderings: one is the inverse of the other.⁷

A positive result of this assumption is that it explains why sentences like (24) are valid.

(24) Jason’s Honda is more dangerous than my Volvo if and only if my Volvo is safer than Jason’s Honda.

If the domains of *safe* and *dangerous* are identical except for the ordering on the objects they contain, and if the ordering of one is the inverse of the other, then any partitioning of the domain of *dangerous* that satisfies the truth conditions of the first conjunct in (24)—i.e., any partitioning that makes ‘*Jason’s Honda is dangerous*’ true and ‘*my Volvo is dangerous*’ false—will have the opposite effect on the domain of *safe*, since the two sets, in effect, stand in the dual relation to each other. For example, a function that partitions the domain of *dangerous* as in (25) must induce a corresponding partitioning on the domain of *safe* as shown in (26), with the result that both conjuncts of (24), shown in (27a)-(27b), are true.

(25) a. $D_{dangerous} = \langle \dots, c, b, my\ Volvo, a, \dots, x, Jason's\ Honda, y, z, \dots \rangle$
b. $pos_d(dangerous) = \langle \dots, x, y, Jason's\ Honda, z, \dots \rangle$
c. $neg_d(dangerous) = \langle a, my\ Volvo, b, c, \dots \rangle$

(26) a. $D_{safe} = \langle \dots, x, y, Jason's\ Honda, z, \dots, a, my\ Volvo, b, c, \dots \rangle$
b. $pos_d(safe) = \langle a, my\ Volvo, b, c, \dots \rangle$

- c. $neg_d(safe) = \langle \dots, x, y, \text{Jason's Honda}, z, \dots \rangle$
- (27) a. $\exists d[(d(dangerous))(Jason's Honda) \& \sim(d(safe))(my Volvo)]$
 b. $\exists d[(d(safe))(my Volvo) \& \sim(d(dangerous))(Jason's Honda)]$

This analysis runs into problems when confronted with examples of CPA, however. Consider (17a), which should have the logical representation in (28).

$$(28) \exists d[(d(tall))(Alice) \& \sim(d(short))(Carmen)]$$

According to (28), (17a) is true just in case there is a function that introduces a partitioning of the domains of *tall* and *short* in such a way that 'Alice is tall' is true and 'Carmen is short' is false; for example, if Alice is very tall and Carmen is not very short. Given the assumption that the domains of the antonymous pair *tall* and *short* have opposite ordering relations, in a context in which the domain of *tall* is (29a), the domain of *short* is (29b).

- (29) a. $D_{tall} = \langle a, b, Carmen, c, Alice \rangle$
 b. $D_{short} = \langle Alice, c, Carmen, b, a \rangle$

In such a context, there is a function that satisfies the truth conditions associated with (28), namely, the one that induces the partitioning of the domains of *tall* and *short* shown in (30) and (31).

- (30) a. $pos_d(tall) = \langle Carmen, c, Alice \rangle$
 b. $neg_d(tall) = \langle a, b \rangle$
 (31) a. $pos_d(short) = \langle b, a \rangle$
 b. $neg_d(short) = \langle Alice, c, Carmen \rangle$

As a result, (17a) should be true. More generally, (17a) should be perfectly well-formed: nothing about the architecture of the analysis predicts that comparatives constructed out of antonymous pairs of adjectives should be anomalous.

The basic problem is that the assumption that the domains of positive and negative adjectives contain the same objects under inverse ordering relations—an assumption that is necessary to account for the validity of sentences like (24)—predicts that it should be possible to interpret sentences like (17a) in the way I have outlined here. One could stipulate that comparison between positive and negative pairs of adjectives is impossible, but there is no aspect of the analysis of comparatives within the partial function approach that derives this constraint. Moreover, such a stipulation would be empirically unmotivated, since comparison of deviation constructions show that comparison between positive and negative adjectives is possible in certain circumstances.

In fact, the interpretations of comparison of deviation constructions provide a second argument against the partial function analysis. On the surface, it appears that a logical representation along the lines of (33b) would actually be an appropriate interpretation for a comparison of deviation construction such as (33a).

- (33) a. The Red Sox are more legitimate than the Orioles are fraudulent.

- b. $\exists d[(d(\textit{legitimate}))(\textit{Red Sox}) \ \& \ \sim(d(\textit{fraudulent}))(\textit{Orioles})]$

As noted above, (33b) would be true in a context in which e.g. The Red Sox are very legitimate and the Orioles are not very fraudulent, which is a possible paraphrase of (33a). The problem is that although such truth conditions are consistent with the meaning of this sentence, they do not account for its entailments: (33b) requires only that there be some partitioning of the domain in which the Red Sox are very legitimate and the Orioles is not very fraudulent, but this partitioning need not be the one associated with the context of utterance. However, as observed in section 3.1, a characteristic of comparison of deviation constructions is that they entail that the properties predicated of the compared objects hold in the context of utterance. If (33b) were the interpretation of (33a), then, this inference would remain unexplained.

3.3 Cross-polar Anomaly and the Ontology of Degrees

A solution to the problem of cross-polar anomaly and comparison of deviation is presented in Kennedy 1997, and worked out in more detail in Kennedy 1999b. (I limit myself here to an outline of the proposal, referring the reader to the above references for more extensive discussion and argumentation.) The solution relies crucially on a characterization of adjectival polarity that is available only in a model in which gradable adjectives map their arguments onto abstract representations of measurement–scales and degrees.

The analysis builds on the intuition that antonymous pairs of adjectives such as *bright~dim* and *tall~short* provide fundamentally the same kind of information about the degree to which an object possesses some gradable property (for, example, both *tall* and *short* provide information about an object’s height), but they do so from complementary perspectives. The positive adjective *tall* is used either neutrally or to highlight the height an object has, while the negative adjective *short* is used to highlight the height an object does not have. This difference in perspective can be exploited in a theory of adjectival polarity in which positive and negative degrees are treated as distinct sorts of objects, an approach that was first suggested by Seuren (1978), and has since been further developed in von Stechow 1984a, Löbner 1990, and Kennedy 1997, 1999a,b (cf. Bierwisch 1989).

The basics of the approach are as follows. A scale *S* can be defined as a linearly ordered, infinite set of points, associated with a dimension that indicates the type of measurement that the scale represents (e.g., *height*, *length*, *weight*, *brightness* and so forth). A degree *d* can then be defined as a convex, nonempty subset of a scale, i.e., a subset of the scale with the following property: $\forall p_1, p_2 \in d \ \forall p_3 \in S [p_1 > p_3 > p_2 \rightarrow p_3 \in d]$ (cf. Landman (1991:110); this is simply the definition of an interval for a linearly ordered set of points).

Assuming that gradable adjectives denote functions from objects to degrees on a scale (Bartsch and Vennemann 1973, Kennedy 1999a), adjectival polarity can be characterized in terms of the sort of degree onto which a particular gradable adjective maps its argument: positive adjectives denote functions from objects to *positive degrees*; negative adjectives denote functions from objects to *negative degrees*. Roughly speaking, positive degrees are intervals that range from the lower end of a scale to some point, and negative degrees are intervals that range from some point to the upper end of a scale. The set of positive and

negative degrees for any scale S ($POS(S)$ and $NEG(S)$, respectively), can be precisely defined as in (34).

- (34) a. $POS(S) = \{d \subseteq S \mid \exists p_1 \in d \forall p_2 \in S [p_2 \leq p_1 \rightarrow p_2 \in d]\}$
 b. $NEG(S) = \{d \subseteq S \mid \exists p_1 \in d \forall p_2 \in S [p_1 \leq p_2 \rightarrow p_2 \in d]\}$

Finally, we can assume that for any object x , the positive and negative projection of x on a scale S ($pos(x)$ and $neg(x)$, respectively) are related as in (35), where MAX and MIN return the maximal and minimal elements of an ordered set.

- (35) $MAX(pos(x)) = MIN(neg(x))$

The result of these assumptions is that the positive and negative projections of an objects x on a scale S are (join) complementary intervals on the scale, as illustrated by the diagram in (36).

- (36) $S: 0 \text{ ————— } pos(x) \text{ ————— } \bullet \text{ ————— } neg(x) \text{ ————— } \rightarrow \infty$

Antonymy, in this view, holds when two adjectives have the same domains but different ranges, and they map identical arguments onto (join) complementary regions of the same scale.

The intuition that the structural distinction between positive and negative degrees is designed to capture is exactly the one I mentioned above: that antonymous pairs of adjectives provide complementary perspectives on the projection of an object onto the scale. This structural distinction is at the core of the account of adjectival polarity outlined here, but more importantly, it also provides the basis for an explanation of cross-polar anomaly in terms of very general principles of ordering relations. A fundamental property of an ordering relation is that its arguments must be elements of the same ordered set; if this requirement is not met, the relation is undefined for the two arguments, and a truth value cannot be computed. A consequence of the analysis of adjectival polarity presented above is that positive and negative adjectives denote functions with different, in fact *disjoint*, ranges: the structural distinction between positive and negative degrees has the consequence that for any scale S , $POS(S)$ and $NEG(S)$ are disjoint. It follows that ordering relations between positive and negative degrees are undefined.

This is the essence of cross-polar anomaly. Assuming that the truth conditions for comparatives are formulated in terms of ordering relations between degrees (the standard assumption in degree-based analyses), any comparative constructed out of adjectives of opposite polarity should fail to have a truth value. For illustration of the analysis, consider the following examples. The logical representation of (37a) in which the adjective in the main clause is negative and the adjective in the comparative clause is positive, is (37b).

- (37) a. Alice is shorter than Carmen is tall.
 b. $short(a) > tall(c)$

The problem is that $short(a)$ and $tall(c)$ denote degrees in different ordered sets: $NEG(height)$ and $POS(height)$, respectively. As a result, the ordering relation introduced by the comparative morpheme is undefined for its two arguments, and the sentence is correctly predicted to be anomalous. (Examples in which the adjectives are reversed are explained in exactly the same way.)

The semantic properties of comparison of deviation constructions can also be explained in this type of model. Since degrees are defined set-theoretically, they are subject to operations on sets. In particular, it is possible to define a difference operation between degrees that returns the amount to which one degree exceeds another one, and the semantics of COD constructions can be defined in terms of orderings between such ‘differential degrees’. Specifically, COD involves a comparison between the degrees to which the compared objects exceed their respective contextually-determined standard values. The logical representation of (38a), for example, is (38b).

- (38) a. The Red Sox are more legitimate than the Orioles are fraudulent.
 b. $(\textit{legitimate}(\textit{Red Sox}) - d_{s(\textit{legitimate})}) > (\textit{fraudulent}(\textit{Orioles}) - d_{s(\textit{fraudulent})})$

According to (38b), (38a) is true just in case the degree to which the Red Sox exceed a standard of legitimacy is greater than the degree to which the Orioles exceed a standard of fraudulence. As shown in Kennedy 1999b (see also Hellan 1981 and von Stechow 1984a,b), ordering relations between such differential degrees is well-defined regardless of whether they are derived from differences between two positive degrees or from differences between two negative degrees. It follows that COD constructions, unlike examples of cross-polar anomaly, should be perfectly interpretable.

Moreover, the logical representation in (38b) crucially accounts for the entailment patterns observed in COD. The semantics of the difference operation is such that the degree to which the Red Sox are legitimate must exceed the standard value for *legitimate* in the context of utterance (likewise for the Orioles and the standard value for *fraudulent*). Since the truth conditions for the non-comparative state that a sentence of the form *x is φ* is true just in case $\varphi(x)$ is at least as great as the standard for φ (see the discussion in section 1), the truth conditions for the non-comparative are satisfied whenever the truth conditions for the comparison of deviation interpretation are satisfied.

Although the preceding discussion is necessarily superficial, it nevertheless makes an important point. In order to construct an analysis of the sort I have outlined here in the first place, it is necessary to make a structural distinction between positive and negative degrees. In order to make this distinction, however, and to use it in turn as the basis for an analysis of adjectival polarity, it must be the case that scales and degrees are part of the ontology, and that the interpretation of gradable adjectives is stated in terms of such objects. Since alternative analyses that do not make reference to scales and degrees do not provide an explanation for cross-polar anomaly, the success of the proposal I have outlined here in this regard can be taken as support for the general hypothesis that the interpretation of gradable adjectives should be characterized in terms of such abstract representations of measurement.

4. Conclusion

Although an analysis of gradable adjectives in terms of families of (possibly partial) functions from individuals to truth values does a good job of explaining most of their semantic properties, and moreover has the advantage of maintaining a simple ontology, it fails to provide an adequate explanation of the distribution

and interpretation of antonymous adjectives in comparatives. Since a principled explanation of these facts can be constructed within a model that characterizes adjectival polarity in terms of a structural distinction between positive and negative degrees (formalized as complementary intervals on a scale), the conclusion to be drawn is that an empirically and explanatorily adequate semantics of gradable adjectives must introduce abstract representations of measurement–degrees qua intervals–into the ontology.

NOTES

* I am very grateful to Chris Barker, Donka Farkas, Bill Ladusaw and Beth Levin for extremely helpful discussion of the material presented here. Errors or inconsistencies in the text are my responsibility.

¹ The standard assumption is that the standard value is set indexically, with respect to some contextually relevant set of objects (a comparison class in Klein’s (1980) terms) that provides the basis for identification of a “norm”, although there are a number of problems with this view (see Klein 1980, Ludlow 1989, and Kennedy 1999).

² In a scalar approach, the comparison class is used to determine the value of the standard-denoting degree (Bierwisch 1989, Ludlow 1989, Kennedy 1999).

³ For two objects x, y in the domain of a gradable adjective φ , $x \geq_{\varphi} y$ iff x is at least as great as y with respect to the implicit ordering on the domain.

⁴ I focus here on comparatives with *more* for perspicuity; see Klein 1980 and Larson 1988 for discussion of equatives and comparatives with *less*.

⁵ The classification of gradable adjectives as positive or negative can be made based on a number of empirical characteristics (see Seuren 1978 for general discussion of this issue). For example, negative adjectives license downward entailments and negative polarity items in clausal complements, but positive adjectives do not (see Seuren 1978, Ladusaw 1979, Linebarger 1980, Sánchez-Valencia 1994, Kennedy to appear); and positive but not negative adjectives can appear with measure phrases (compare ‘2 meters long’ with ‘?2 meters short’).

⁶ A third set of facts that bear on the analysis of CPA in particular and antonymy more generally, but which I will not consider in this paper, is illustrated by the sentences in (i). (I am grateful to Chris Barker (personal communication) for first bringing these facts to my attention.)

- (i) a. The C is sharper than the D is flat.
- b. My watch is faster than your watch is slow.
- c. She was earlier than I was late.

While these sentences appear to counterexemplify the generalization adduced from CPA—that antonymous adjectives are anomalous in comparatives—the situation is in fact more complicated. As shown in Kennedy 1997, 1999a,b, both members of the ‘antonymous’ pairs in these constructions have the properties of *positive* adjectives.

⁷ This idea is implicit in Klein's (1980:35) discussion of examples like *Mona is more happy than Jude is sad* (see the discussion of comparison of deviation below). Sánchez-Valencia (1994) shows how this assumption can be used to build an explanation of the monotonicity properties of polar adjectives.

REFERENCES

- Atlas, J. 1984. Comparative adjectives and adverbials of degree: An introduction to radically radical pragmatics. *Linguistics & Philosophy* 7:347-377.
- Bartsch, R. and T. Vennemann. 1973. *Semantic Structures: A Study in the Relation between Syntax and Semantics*. Frankfurt: Athenaeum Verlag.
- Bierwisch, M. 1989. The semantics of gradation. In M. Bierwisch and E. Lang (eds.), *Dimensional Adjectives*. Berlin: Springer-Verlag. 71-262.
- Cresswell, M.J. 1976. The semantics of degree. In B. Partee (ed.), *Montague Grammar*. New York: Academic Press. 261-292.
- Hale, A. 1970. Conditions on English comparative clause pairings. In R. Jacobs and P. Rosenbaum (eds.) *Readings in English Transformational Grammar*. Washington: Georgetown University Press.
- Heim, I. 1985. Notes on comparatives and related matters. Ms., University of Texas, Austin.
- Hellan, L. 1981. *Towards an Integrated Analysis of Comparatives*. Tübingen: Narr.
- Hendriks, P. 1995. *Comparatives and Categorical Grammar*. Ph.D. thesis, University of Groningen.
- Hoeksema, J. 1983. Negative polarity and the comparative. *Natural Language and Linguistic Theory* 1:403-434.
- Kamp, H. 1975. Two theories of Adjectives. In Keenan, Edward. (ed.) *Formal Semantics of Natural Language*. Cambridge: Cambridge University Press, 123-155.
- Kennedy, C. 1999a. *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. New York: Garland.
- Kennedy, C. 1999b. Polar opposition and the ontology of "degrees". Unpublished ms., Northwestern University.
- Kennedy, C. To appear. On the monotonicity of polar adjectives. In J. Hoeksema, H. Rullmann, V. Sánchez-Valencia and T. van der Wouden (eds.), *Perspectives on Negation and Polarity*. Amsterdam: Benjamins.
- Klein, E. 1991. Comparatives. In von Stechow, A. & D. Wunderlich (eds.) *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*. Berlin: Walter de Gruyter.
- Klein, E. 1982. The interpretation of adjectival comparatives. *Journal of Linguistics* 18:113-136.
- Klein, E. 1980. A semantics for positive and comparative adjectives. *Linguistics & Philosophy* 4.1:1-45.
- Ladusaw, William. 1979. *Polarity as Inherent Scope Relations*. Ph.D. dissertation, University of Texas, Austin.
- Landman, F. 1991. *Structures for semantics*. Dordrecht: Kluwer.
- Larson, R. 1988. Scope and comparatives. *Linguistics & Philosophy* 11:1-26.
- Linebarger, 1980. *The Grammar of Negative Polarity*. Ph.D. dissertation, MIT.
- Löbner, S. 1990. *Wahr neben Falsch: Duale Operatoren als die Quantoren natürlicher Sprache*. Tübingen: Niemeyer.
- McConnell-Ginet, S. 1973. *Comparative Constructions in English: A Syntactic and Semantic Analysis*. Doctoral dissertation, University of Rochester.

- Moltmann, F. 1992. *Coordination and Comparatives*. Ph.D. thesis, Massachusetts Institute of Technology.
- Rullmann, H. 1995. *Maximality in the Semantics of WH-Constructions*. Ph.D. thesis, UMass., Amherst.
- Sánchez-Valencia, V. 1994. Polarity, predicates, and monotonicity. Ms., University of Groningen.
- Sapir, Edward. 1944. Grading: A study in semantics. *Philosophy of Science* 11:93-116.
- Seuren, P.A.M. 1984. The comparative revisited. *Journal of Semantics* 3:109-141.
- Seuren, P.A.M. 1978. The structure and selection of positive and negative gradable adjectives. In Farkas, D., W. Jacobson & K Todrys (eds.), *Papers from the Parasession on the Lexicon*, CLS 14, University of Chicago. 336-346.
- Seuren, P.A.M. 1973. The comparative. In F. Kiefer & N. Ruwet (eds.), *Generative Grammar in Europe*. Dordrecht: Riedel.
- von Stechow, A. 1984a. Comparing semantic theories of comparison. *Journal of Semantics* 3:1 77.
- von Stechow, A. 1984b. My reaction to Cresswell's, Hellan's, Hoeksema's, and Seuren's comments. *Journal of Semantics* 3:183-199.