Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Predicates *

.

Christopher Kennedy (ck@uchicago.edu)
University of Chicago

DRAFT: April 5, 2006 (comments welcome)

Abstract. This paper investigates the way that grammatical (lexical semantic) features of linguistic expressions influence vagueness, focusing on the interpretation of the positive (unmarked) form of gradable adjectives. I begin by developing a semantic analysis of the positive form of 'relative' gradable adjectives, in which vagueness derives from the truth conditions of the predicate, which require an object to possess a contextually significant degree of the relevant property (as in Graff 2000). The analysis expands on previous proposals both in further motivating a semantic approach to the vagueness of the positive form and in precisely identifying and characterizing the division of labor between compositional and context dependent elements in its interpretation. I then introduce a challenge to the analysis from the class of 'absolute' gradable adjectives: adjectives that are demonstrably gradable, but which have positive forms that relate objects to maximal or minimal degrees, and do not give rise to vagueness. I argue that the truth conditional difference between relative and absolute adjectives in the positive form stems from the interaction of lexical semantic properties of gradable adjectives and a general constraint on interpretive economy that requires the meaning of a constituent to be computed strictly on the basis of the conventional meanings of its subconstituents to the extent possible, allowing for context dependent truth conditions only as a last resort.

Keywords: vagueness, context dependence, borderline cases, Sorites Paradox, gradable adjectives, standard of comparison, comparison class, scale structure

1. Introduction

The general question that this paper addresses is how precisely sentences like (1) are assigned truth conditions in a context of utterance.

(1) The coffee in Rome is expensive.

The problem presented by sentences of this sort is that they are vague: what exactly it means to 'count as' expensive is indeterminate. Sentences like (1) have three distinguishing characteristics, which have been the focus of much work on vagueness in semantics and the philosophy of language. The first is truth conditional variability: (1) could be judged true if asserted as part of a conversation about the cost of

* To appear in *Linguistics and Philosophy*

living in various Italian cities (*In Rome, even the coffee is expensive!*), for example, but false in a discussion of the cost of living in Chicago vs. Rome (*The rents are high in Rome, but at least the coffee is not expensive!*)

The second feature of vagueness is the existence of 'borderline cases'. For any context, in addition to the sets of objects that a predicate like *is expensive* is clearly true of and clearly false of, there is typically a third set of objects for which it is difficult or impossible to make these judgments. Just as it is easy to imagine contexts in which (1) is clearly true and contexts in which it is clearly false, it is also easy to imagine a context in which such a decision cannot be so easily made. Consider, for example, a visit to a coffee shop to buy a pound of coffee. The Mud Blend at $1.50/pound is clearly not expensive, and the Organic Kona at $20/pound is clearly expensive, but what about the Rocket to the Moon Blend at $9.25/pound? A natural response is 'it depends' (on the price of other blends, on how much I am willing to spend, etc.); this is the essence of being a borderline case.

Finally, vague predicates give rise to the well-known Sorites Paradox when inserted in arguments with the form shown in (2).

(2) *The Sorites Paradox*

P1. A $5 cup of coffee is expensive (for a cup of coffee).

P2. Any cup of coffee that costs 1 cent less than an expensive one is expensive (for a cup of coffee).

C. Therefore, any free cup of coffee is expensive.

The structure of the argument appears to be valid, and the premises appear to be true, but the conclusion is without a doubt false. It is clear is that the problem with the argument lies somewhere in the induction built on the second premise; what is hard is figuring out exactly what goes wrong.

It is widely accepted that the locus of vagueness in sentences like (1) is the predicate headed by the gradable adjective *expensive* (though vagueness is not restricted to gradable adjectives, a point I will return to later). Within linguistic semantics, a fruitful line of research has developed that analyzes an unmarked gradable predicate — henceforth the POSITIVE FORM of the predicate — as a relation between the degree to which an object posessess the property expressed by the predicate and a context dependent STANDARD OF COMPARISON for the property. A predicate like *is expensive* denotes the property of having a degree of cost that is at least as great as some standard of comparison of cost, where the value of the standard is crucially not part of the the lexical meaning of *expensive*, but is rather determined contextually.

Truth-conditional variability arises when the standard of comparison is shifted: if the standard of comparison for *expensive* based on the cost of coffee in Italian cities is lower than the cost of coffee in Rome, (1) true; if the standard based on the cost of coffee in Rome and Chicago is higher than the cost of coffee in Rome, (1) is false.

There is general agreement among researchers in linguistics that something like this is what is going on in the interpretation of gradable predicates, but several fundamental questions remain open. The first is the question of how the semantic analysis of gradable predicates relates to the phenomena of borderline cases and the Sorites Paradox. Specifically: what feature of the semantics of gradable predicates is responsible for their behavior with respect to these two (more general) characteristics of vague expressions? The second is the question of what the actual content of the standard of comparison is and how it is computed. In particular, to what extent is its value explicitly encoded in the semantic representation of gradable predicates and/or determined by the conventional meanings of various subconstituents of such predicates, and to what extent is it determined by purely contextual, possibly extra-linguistic factors? These questions have received a fair amount discussion in the literature (see e.g., Sapir 1944; Lewis 1970; Wheeler 1972; McConnell-Ginet 1973; Kamp 1975; Fine 1975; Klein 1980; Pinkal 1995; von Stechow 1984; Bierwisch 1989; Ludlow 1989; Kennedy 1999; Barker 2002), but a fully comprehensive theory has not been developed, mainly because the full range of relevant data has not been taken into account.

In particular, most analyses fail to address the distinction between RELATIVE gradable adjectives like *expensive*, which have the properties of vagueness described above, and ABSOLUTE gradable adjectives like *straight* and *bent*, which do not. As I will show in detail, predicates like *straight* (as in *The rod is straight*) require their arguments to posess a maximal degree of the relevant property, and those like *bent* (as in *The rod is bent*) merely require their arguments to posess a non-zero degree of the relevant property; neither describes a relation to a context dependent standard of comparison. However, despite these differences in interpretation, relative and absolute gradable adjectives are the same semantic type, and express the same kind of meanings. We should therefore expect an explanatorily adequate theory of the positive form to derive their differences, not stipulate them.

The goal of this paper is to develop such a theory, and in doing so, to answer the two questions outlined above. I will begin by developing a semantics for the positive form of relative adjectives that improves in several ways on previous analyses, both in the way that it explains the vagueness of this form, and in the way it accounts for the division of

labor between the compositional and contextual elements of its interpretation. I will then introduce the problem for the analysis presented by absolute adjectives and provide detailed empirical evidence for the relative/absolute distinction. Finally, I will propose an analysis of the relative/absolute distinction in which the positive form is polysemous, and the interpretation of predicates headed by relative and absolute adjectives is determined by their lexical semantic properties, in particular the structures of the scales that represent the type of gradable properties they encode, and a general constraint on 'interpretive economy' that requires the denotation of a constituent to be computed on the basis of conventional properties of its subconstituents to the extent possible, allowing for context sensitive features of meaning only as a last resort. I conclude by discussing the implications of the facts anlyzed here for alternative semantic analyses of vagueness in gradable predicates, focusing in particular on approaches that eschew a degree-based semantics in favor of a more general (and arguably independently necessary) supervaluation semantics.

2. The semantics of the positive form

2.1. Gradable adjectives and degree morphology

I begin with an overview of the semantic analysis of gradable adjectives and the constructions in which they appear. My core assumptions about gradable adjective meaning, which is shared in some form by most other analyses (see e.g., Seuren 1973; Bartsch and Vennemann 1973; Cresswell 1977; Hellan 1981; von Stechow 1984; Heim 1985, 2000; Bierwisch 1989; Klein 1991; Kennedy 1999, Kennedy and McNally 2005), are stated in (3).

(3) a. Gradable adjectives map their arguments onto abstract representations of measurement, or DEGREES.
b. A set of degrees totally ordered with respect to some DIMENSION (height, cost, etc.) constitutes a SCALE.

In other words, I assume a semantic ontology that includes the type 'degree' (d) along with individuals, truth values, possible worlds, and so forth.[1]

[1] The leading contenders to scalar analyses of gradable adjective meaning are those that treat gradable adjectives as partial functions from individuals to truth values with context dependent extensions and adopt a supervaluation analysis of vagueness (see e.g., McConnell-Ginet 1973, Kamp 1975, Fine 1975, Klein 1980, Pinkal 1995). I will discuss these sorts of approaches in section 4.4, where I will argue

There are various compositional implementations of the core hypotheses about gradable adjective meaning stated in (3); here I will follow Bartsch and Vennemann 1972, 1973 and Kennedy 1999 and analyze gradable adjectives as measure functions: functions of type $\langle e, d\rangle$. The adjective *expensive*, for example, is a function from the subset of the domain of individuals that have some cost value to (positive) degrees of cost. I will represent gradable adjective denotations as in (4), where $\mathbf{adj}(x)$ should be understood as an abbreviation for 'the degree on the appropriate scale that represents x's measure of adjective-ness'.[2]

(4) a. $[\![\text{expensive}]\!] = \lambda x.\mathbf{expensive}(x)$
b. $[\![\text{tall}]\!] = \lambda x.\mathbf{tall}(x)$
c. $[\![\text{old}]\!] = \lambda x.\mathbf{old}(x)$

Measure functions are converted into properties of individals by degree morphology, which in English includes (at least) the comparative morphemes (*more, less, as*), intensifiers (*very, quite, rather*, etc.), the sufficiency morphemes (*too, enough, so*), the question word *how*, and so forth. Degree morphemes serve two semantic functions: they introduce an individual argument for the measure function denoted by the adjective, and they impose some requirement on the degree derived by applying the adjective to its argument, typically by relating it to another degree. Syntactically, I assume that gradable adjectives project extended functional structure headed by degree morphology (Abney 1987,Corver 1990,Grimshaw 1991, Kennedy 1999), and that the adjectival projection is thus a Degree Phrase, rather than an Adjective Phrase.

As an illustration, consider the structure and interpretation of the comparative predicate in (5a), shown in (5b).

(5) a. Chicago is larger than Rome.

---

that even if we determine that supervaluations are the best way to handle various types of semantic imprecision (including vagueness), we still need to characterize the meanings of gradable predicates in terms of scales and degrees, since these are the features that crucially explain the semantic differences between the classes of gradable adjectives discussed in this paper.

[2] A note on notation: throughout this paper I follow Heim and Kratzer 1998 in assuming that syntactic representations can be directly interpreted, but I will use predicate logic as my metalanguage for representing truth conditions, rather than English as in Heim and Kratzer 1998, defining any new symbols that I introduce (as for **expensive**, etc. above). I will use lambda notation to represent functional meanings, and to keep the representations as simple as possible, I will omit type specifications for arguments from the domain of individuals, degrees and gradable adjectives, instead using the variables from the sets $\{x, y, z\}, \{d, d', d''\}, \{g, g', g''\}$ respectively. Finally, I will omit specification of assignment functions and other contextual parameters except where relevant.

b.

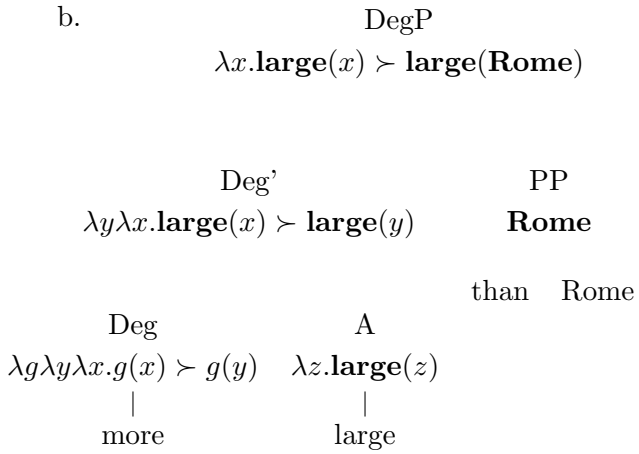

Here *more* is treated as an expression that establishes an ordering relation between two degrees: one derived by applying the adjectival head to its subject, the other by applying it to the 'standard' constituent, marked by *than* (Hankamer 1973; Hoeksema 1984 Heim 1985; Kennedy 1999).[3] Composition derives the property at the top of (5b) as the denotation of the DegP: the comparative predicate *larger than Rome* is true of an object if the degree to which it is large exceeds the degree to which Rome is large.[4]

[3] I give a 'phrasal' semantics for *more* here for simplicity; an alternative is that the standard constituent is an elided clause that directly denotes a degree (see the references cited above for comparison of the two approaches). I also assume that the relation between *more large* and *larger* is a matter of phonology.

[4] This example can be used to illustrate a phenomenon that often arises in discussions of vagueness, but which I will not address in detail here: INDETERMINACY (McConnell-Ginet 1973; Kamp 1975; Klein 1980; Kennedy 1999). Indeterminacy is the possibility of associating a single lexical item with several distinct but related measure functions. *Large* in (5a), for example, can be used either to measure population or to measure sprawl, resulting in distinct truth conditions. (For example, if the population of Rome were doubled, (5a) would be false on the population reading but would remain true on the sprawl reading.) More complex cases are adjectives like *skillful* and *clever*, which are highly underspecified for the precise feature being measured. Although both indeterminacy and vagueness are factors that need to be resolved in order to derive determinate truth conditions for sentences constructed out of gradable adjectives, I assume that the former is distinct from the latter (Pinkal 1995; Kennedy 1999). In particular indeterminacy is a feature of adjectives generally (it is a kind of polysemy), while vagueness is a feature of the positive form specifically, as shown by the fact that (5a) is indeterminate but not vague. Indeterminacy and vagueness do interact, however, since the resolution of the former is a prerequisite for the resolution of the latter. This follows from the analysis to be developed in this paper, in which the computation of the standard of comparison in the positive form is dependent on fixing a denotation for the adjective.

This sort of analysis can in principle be extended to any other degree morpheme/adjective combination — collocations with *less, as, too, enough, so, how, very, quite, rather* and so forth — with appropriate modifications to the denotations of the degree morphemes. Whether these are ultimately the correct assumptions about the semantic type of gradable adjectives and the syntax of the adjectival projection is a question that I will not address here, since the answer depends factors that are independent of the general issue addressed in this paper: the interpretation of the positive form.[5]

2.2. The positive form

Paradoxically, it is the most morphosyntactically simple form of a gradable predicate — the unmarked, positive form — that creates the most problems for a compositional semantic analysis.[6] The first (and easiest) problem to overcome is a morphological/type theoretic one: if gradable adjectives denote functions of type $\langle e, d\rangle$, then they must combine with degree morphology to derive a property of individuals.

[5] The main alternative analysis is one in which gradable adjectives denote relations between degrees and individuals (they are type $\langle d, \langle e, t\rangle\rangle$), and comparatives and other degree constructions saturate the degree argument of the predicate; see Kennedy 1999; Heim 2000; Meier 2003; Bhatt and Pancheva 2004; Neeleman et al. 2004 for relevant discussion of the issues at stake in choosing between the two approaches. I adopt the measure function analysis of gradable adjectives and degree morphology primarily because it provides a transparent interpretation scheme for the 'extended projection' syntax for the adjectival projection illustrated in (5b), which has a wide range of empirical and theoretical support, and because it makes it easy to see which bits of structure are contributing which bits of meaning. However, since the relational analysis also assumes that gradable adjecives encode measure functions as part of their meaings (for example, *tall* holds of an individual a and a degree d just in case a's height is at least as great as d), all of my central proposals could be adapted to this type of approach with appropriate changes in semantic type and denotation of the relevant constituents. The crucial assumptions are the ones stated in (3), which are shared by all scalar analyses of gradable predicates.

[6] I distinguish the completely unmarked form of e.g. *old* in (6c) from the form in (i), which is not combined with a bound degree morpheme, but is combined with a measure phrase.

(i) That dog is 2 years old.

The semantic justification for this distinction is that the two forms have quite different truth conditions: (6c) establishes a relation between the subject's age and a contextual standard of comparison (in a way that will be described in detail below), while (i) relates the subject's age to whatever arbitrary degree is denoted by the measure phrase. In terms of compositional semantics, we may assume either that measure phrases combine directly with gradable adjectives (Klein 1980; Kennedy and McNally 2005), or that a specialized Deg head mediates the relation between them (Svenonious and Kennedy 2005; cf. Schwarzschild 2004).

The positive form does not have any overt degree morphology, however, as illustrated by examples like (6).

(6) a. The Mars Pathfinder mission was expensive.
b. My six-month old son is big.
c. That dog is old.

One solution to this problem would be to take the absence of morphology at face value and assume that the positive form is a simple AP (i.e., an adjectival projection without functional degree morphology), but that the grammar includes a type-shifting rule that turns measure functions into properties of individuals (see e.g. Neeleman et al. 2004). A second solution would be to assume that the DegP in the positive form is headed by a null morpheme that has the same semantic function as overt degree morphology: it takes a gradable adjective denotation (a measure function) and returns a property of individuals (see e.g., Bartsch and Vennemann 1972; Cresswell 1977; von Stechow 1984; Kennedy 1999).

Potential evidence in favor of such a morpheme comes from Mandarin Chinese. As shown by Sybesma (1999), the positive form of the adjective in Mandarin is morphologically marked by the morpheme *hen*, as shown in (7a). *Hen* is sometimes glossed as *very*, but it also has a neutral interpretation that just marks the positive form (see Sybesma 1999, p. 27 for discussion). The unmarked form in (7b) is infelicitous if uttered in isolation, but in a context that supports a comparative interpretation, it is acceptable.

(7) a. Zhangsan hen gao.
Zhangsan HEN tall
'Zhangsan is tall.'
b. Zhangsan gao.
Zhangsan tall
'Zhangsan is taller (than X).'

Furthermore, as shown by the examples in (8a-b), fully explicit comparative constructions are incompatible with *hen*. This follows if the comparative degree morpheme(s) in this language are null heads (of category Deg), and so in complementary distribution with *hen*.

(8) a. Zhangsan bi ni (*hen) gao.
Zhangsan than you (*HEN) tall
'Zhangsan is taller than you.'
b. Zhangsan (*hen) gao-de neg mozhao tianpeng.
Zhangsan (*HEN) tall-DE can touch ceiling
'Zhangsan is so tall that he can touch the ceiling'

These facts suggest that *hen* is a morphological realization of a 'positive degree morpheme'. The Mandarin data do not provide conclusive evidence that English also has such a morpheme — it may be the case that Mandarin has lexicalized a meaning that in English is encoded as a lexical type-shifting rule — but it is certainly compatible with this view. In order to develop a compositional analysis of the positive form that is fully paralel to the analysis of forms with overt degree morphology, I will assume in this paper that English has a phonologically null version of Mandarin *hen*, which I will refer to as *pos* (for 'positive form'). However, nothing crucial hinges on this assumption: the content of my proposals and argumentation remains the same if we assume instead that 'the positive degree morpheme *pos*' is really 'the positive type-shifting rule *pos*'.

The harder question to answer is the semantic one: what is the meaning of *pos*? If we accept the assumptions about gradable adjective meaning outlined in the previous section, then an answer to this question constitutes an answer to the questions I began this paper with. That is, if gradable adjectives themselves have fixed denotations as measure functions (modulo indeterminacy; see note 4), then any characteristics of vagueness associated with the positive form that stem from aspects of conventional linguistic meaning must be located in the semantics of *pos* and its interaction with other constituents of the sentence.

In section 1, I characterized the denotation of the positive form of *expensive* in terms of a relation to a contextually-determined standard of comparison: *is expensive* denotes the property of having a degree of cost that is at least as great as the prevailing standard. This means that *pos* should be assigned a denotation along the lines of (9), where $\mathbf{d}_s$ is shorthand for 'the contextually appropriate standard of comparison, whatever that is'.

(9) $\quad [\![[_{\text{Deg}}\ pos]]\!] = \lambda g \lambda x.g(x) \succeq \mathbf{d}_s$

Now we can rephrase the central question in an even more precise way: is the value of the standard of comparison compositionally determined in a way specified by the conventional meaning of *pos*, or is the standard merely a variable introduced by *pos*, whose value is a function of extralinguistic factors?

The main argument in favor of the former view is that the standard of comparison can be manipulated in what appears to be a compositional way by constituents local to the predicate, such as the *for*-PPs and modified nominals in (10a-c) and (11a-c).

(10) a. Kyle's car is expensive for a Honda.

b. Nadia is tall for a gymnast.
c. Jumbo is small for an elephant.

(11) a. Kyle's car is an expensive Honda.
b. Nadia is a tall gymnast.
c. Jumbo is a small elephant.

These constituents clearly have an effect on truth conditions: (10c) and (11c) could both be true in a situation in which the 'bare' positive in (12a) is false, as illustrated by the fact that (12b-c) are non-contradictory.

(12) a. Jumbo is small.
b. Jumbo is small for an elephant, but he is not small.
c. Jumbo is a small elephant, but he is not small.

A common interpretation of these facts is that the standard of comparison is always computed relative to a COMPARISON CLASS, which can be made explicit by a *for*-PP or modified nominal (Klein 1980). It may also remain implict, however. Taking into account the fact that bare positives like (12a) can be understood as equivalent to their variants in examples like those in (10) and (11) (in appropriate contexts), many analyses assume that the comparison class is always a constituent of the semantic representation of the positive form (Bartsch and Vennemann 1972; Wheeler 1972; Cresswell 1977; von Stechow 1984; Kennedy 1999; Kennedy and McNally 2005).

Bartsch and Vennemann (1972), for example, provide a denotation for *pos* that is essentially equivalent to (13), where c is a property and **norm** is a function that returns the average degree to which the objects in the set defined by c (the comparison class) possess the gradable property measured by g.

(13) $[\![[_{\text{Deg}}\ pos]]\!] = \lambda g \lambda c \in D_{\langle e,t\rangle} \lambda x.g(x) \succ \mathbf{norm}(c)(g)$

If we assume that in the absence of explicit information about the value of c, this argument can be saturated by a variable over properties (perhaps syntactically encoded as a null constituent; an idea that I elaborate on below), the contextual variability of the positive form boils down to the task of finding an appropriate property for the comparison class variable. This analysis therefore has the advantage resolving the context dependence of the positive form using independently necessary mechanisms for resolving property variables (which are needed to handle e.g. *one*-anaphora and other types of predicate anaphora), rather than by introducing an additional contextual parameter just for the purpose of fixing the standard of comparison of vague predicates.

Further evidence in favor of the analysis of the positive form in (13), and further support for the hypothesis that the standard of comparison is compositionally determined, comes from the fact that the standard can vary as a function of the value of the argument of the predicate. This is illustrated most clearly by an example like (14a), which can have the truth conditions stated in (14b) (Kennedy 1999).

(14) a. Everyone in my family is tall.
b. for every x such that x is a member of my family, x has a height greater than the norm for a comparison class based on x.

What is important about this example is that the standard of comparison can vary with the quantificational subject, indicating that some type binding relation holds between the subject and the comparison class variable (Ludlow 1989, Stanley 2002). If the standard of comparison were simply a free variable over degrees whose value is contextually determined, there would be no way to represent this relation.

The analysis in (13) can handle this data by assuming that an implicit comparison class variable can range not just over properties, but over properties that are relativized to particular individuals, i.e., skolem functions. Using $\mathbf{c}_x$ to represent a variable over properties 'related to x' (the property of being the same age as x, the property of being the same sex as x, etc.), the truth conditions of (14a) can be accurately represented as in (15).

(15) $\forall x[\mathbf{member\text{-}of\text{-}my\text{-}family}(x) \rightarrow \mathbf{tall}(x) \succ \mathbf{norm}(\mathbf{c}_x)(\mathbf{tall})]$

The linguistic evidence for such functions is well-established in work on e.g. functional readings of pronouns, functional questions and choice-function analyses of indefinites, so again nothing particularly new needs to be added to the theory to account for facts like (14a).[7]

2.3. Eliminating norms

Despite these advantages, there are a couple of problems with the analysis of the positive form outlined in the previous section. The first is that reducing the context dependence of the positive form to the identification of a comparison class fails to explain the fact that the positive form of a gradable predicate gives rise to borderline cases and

[7] I leave aside the details of how exactly this sort of analysis would be implemented compositionally, since I will suggest an alternative below. See Ludlow 1989 and Stanley 2002 for proposals.

the Sorites Paradox, and even worse, the semantic analysis stated in (13) derives the wrong truth conditions.

As noted above, an apparent advantage of the comparison class analysis of the positive form as implemented in (13) is that it captures context dependence in terms of general, independently motivated, interpretive mechanisms: the positive form is just another example of a construction that includes an implicit property variable. Once that variable is fixed by the contextual assignment function, the actual value of the standard of comparison can be computed strictly on the basis of the conventional meaning of *pos*: no further role of context is required. In particular, according to (13), the standard of comparison will always be a degree on the scale of the adjective that represents the average degree to which the objects in the comparison class possess the property measured by the adjective.

The problem with this analysis is that it fails to provide a satisfactory explanation of borderline cases or the Sorites Paradox (Rusiecki 1985; Pinkal 1995). Once the comparison class is fixed, then the cutoff point for the objects that the positive form is true of is also fixed: it is the average degree to which the objects in the comparison class possess the property. As far as the truth conditions of the positive form are concerned, then, there should be no borderline cases: if an object possesses the relevant property to a degree at least as great as the average, the positive form is true of it, otherwise false. Likewise, these truth conditions should entail that the second premise of the Sorites argument is false: at some point in a Sorites sequence the average will be crossed, falsifying the universal generalization expressed by the premise.

One possible response to this problem would be to argue that it is not the meaning of the positive form *per se* that is responsible for these features of the positive form, but rather the fact that sentences with implicit comparison classes are consistent with an infinite number of possible interpretations (corresponding to different ways of fixing the comparison class variable). The indeterminacy that arises from this ambiguity in turn gives rise to borderline cases and the Sorites Paradox. However, this explanation would predict that the addition of an explicit comparison class should eliminate borderline cases and the Sorites Paradox, a prediction that is not correct, as pointed out by Graff (2000).

Consider for example (16). The comparison class is explicit, but it is perfectly plausible that one could know that the median rent for apartments on the street is $700 and still be unwilling to judge this sentence as true if, for example, there are a few expensive apartments with rents significantly higher than $725.

(16) A rent of $725 is expensive for an apartment on this street.

Similarly, the argument in (17) remains just as paradoxical as the one discussed in section 1, even though the generalization in (17P2) should fail for the move from $701 to $700.

(17) P1. A rent of $1000 is expensive for an apartment on this street.
P2. A rent that is $1 less than an expensive rent is expensive for an apartment on this street.
C. A rent of $100 is expensive for an apartment on this street.

In short, if the single parameter of contextual variation in the positive form is the comparison class, then the resulting truth conditions are too precise to support a semantic explanation of borderline cases and the Sorites Paradox.

A second respose to this would be to claim that these features of vague predicates don't have a semantic explanation in the first place, and so don't bear on the characterization of the truth conditions of the positive in terms of an average degree (for a comparison class). For example, it could be the case that the truth conditions of the positive form are as in (13), but for purely epistemological reasons, we can never know for sure where the standard is: we know that it is an average for a comparison class, but we can never know its actual value (Williamson 1994). The 'region of uncertainty' surrounding the standard gives rise to borderline cases, and explain why we are unable to pinpoint the precise location in a graded sequence where the second premise of the Sorites Paradox fails.

One objection to this explanation, pointed out by Graff (2000), is that it doesn't explain why we are willing to accept the second premise as true. An even bigger problem, however, is pointed out by Bogusławski (1975): the analysis in (13) does not derive the right truth conditions, and so fails on purely linguistic grounds. If the positive form means 'have a degree of property g greater than the average for a comparison class based on g', then (18) should be a contradiction:

(18) Nadia's height is greater than the average height of a gymnast, but she is still not tall for a gymnast.

(18) is not a contradiction, however, providing clear semantic evidence that standards are not averages. This problem could be fixed by re-defining **norm** so that it identifies some value other than an average for a comparison class, but any such characterization that is not itself context dependent will reproduce the problems described above. The

conclusion, then, is that some element of the truth conditions of the positive form other than just the comparison class must be context dependent. In other words, it is not enough to make the argument of the standard identifying function context dependent; the function itself must be context dependent (in some way) as well.

A solution to both of these problems is offered by Graff (2000) in her analysis of the Sorites Paradox. Graff argues that the positive form does not merely require an object to meet a norm for a comparison class, but rather that an object must exceed the norm to a degree that is *significant* given our interests (cf. Bogusławski 1975). She implements her analysis by introducing a new ordering relation $!\succ$ which holds between two degrees iff the first is 'significantly greater than' the second, and defining *pos* as in (19) (where **norm** is as above).

(19) $\quad [\![[_{\text{Deg}}\ pos]]\!] = \lambda g \lambda c \in D_{\langle e,t\rangle} \lambda x.g(x)\ !\succ \mathbf{norm}(c)(g)$

This analysis clearly explains why (18) is not a contradiction: it is possible to exceed the norm without exceeding it by a significant amount; if the latter relation is not satisfied, the positive form does not hold. More importantly, the analysis provides a semantic account of borderline cases and the Sorites Paradox. The former arise from indeterminacy about what counts as significant, and the latter — as well as our reactions to it — derives from the interest relativity inherent in this notion.

Specifically, according to Graff, we typically have an interest in efficiency, which yields the following result: when x and y are extremely similar with respect to some gradable property, and they are being actively considered, the cost of discriminating between them with respect to that property typically outweighs the benefit. As a result, they count as 'the same for present purposes', and one is significantly more than the norm for the property if and only if the other is. In other words, whenever you look at two objects in a Sorites sequence to see if one is significantly greater than the norm and the other isn't, you raise their similarity enough to ensure that such a result would be impossible. This is why we have the (incorrect) intuition that the second premise of the paradox is true.

In effect, Graff's analysis introduces a context dependent property of degrees into the truth conditions of the positive form: her $!\succ$ relation is equivalent to the standard partial ordering relation $\succeq$ plus an extra requirement that there exist a significant difference between the related degrees, where what counts as significant is a matter of context, subject to the interests and expectations of the participants in a discourse, as well other contextual factors. It is therefore worth asking whether

the truth conditions of the postive form make reference to a norm or average at all: why not instead assume that the positive simply requires its argument to possess a significant degree of the relevant property (based on a comparison class), and dispense with norms and averages completely? This is essentially the proposal made by Bogusławski (1975), who argues that the positive form requires an object to possess the relevant property to a degree that is 'conspicuous', or 'sufficient to attract attention'.

Evidence in favor of this hypothesis comes from the fact that in certain contexts, the standard of comparison can actually be less than an average degree, as Graff herself observes. For example, if I walk into a bar that I know to have extremely cheap but delicious coffee, and find that they have suddenly raised the price of an espresso, I can felicitiously use (20) to complain about this to the baristo even if the actual cost of the coffee remains below average.

(20) Hey, the coffee is expensive now!

In this example, the new price is significant relative to my expectations, but still below average.

Given these considerations, I will adopt a version of Graff's proposal in which the standard-identifying function is itself context dependent, and the meaning of the positive form similar to that suggested by Boguslawski: *is expensive* is true of an object if the degree to which it is expensive is 'significant'. Typcially, only a larger than average degree of a particular property will count as significant because part of this notion involves differentiating an object from other objects with respect to the degree to which they possess some property. Assuming that if an object a possesses property g to degree d, then it possess g to any degree less than d, smaller degrees will typically fail to provide meaningful differentiations. However, in contexts in which some other feature provides a basis for determining significance, such as the expectations of the participants in the discourse, this need not be the case.

Before presenting a formal implementation of the revised proposal, however, I need to address the second problem with the 'standard' semantics of the positive form, which has to do with the status of the comparison class as a semantic argument.

2.4. Eliminating comparison classes

The evidence that comparison classes are semantic arguments came from *for*-PPs and modified nominals, which appear to affect the computation of the standard of comparison in a compositional way, and

from the fact that they can be bound by quantifiers external to the predicate. On closer inspection, however, it turns out that none of these facts support this conlcusion.

Consider first the case of modified nominals. Although such nominals typically provide the basis for computing the standard of comparison for adjectival predicates that modify them, they do not have to. As shown by (21a), it is possible to assert that something is *an A NP* while denying that it is *A for an NP*.[8]

(21) Kyle's car is an expensive BMW, though it's not expensive for a BMW. In fact, it's the least expensive model they make.

The conclusion to draw from (21) is that although the denotation of the modified nominal provides a highly salient property which may be used to calcluate the standard of comparison, there is nothing compositional about this relation, and nothing obligatory (contra Wheeler 1972). Consequently, there is no reason to conclude based on the tendency for a modified nominal to provide the comparison class that the latter is a constituent of the conventional meaning of the positive form, as opposed to just some other bit of relevant contextual information: the facts could be handled just as well by an analysis in which the standard of comparison is a variable whose value is conventionally determined on the basis of some discourse-salient property (cf. Pinkal 1979). A modified noun denotation is arguably the most salient property at the point of interpreting the adjectival predicate, explaining the strong tendency for it to be used as the comparison class, but facts like (21) show that it does not have to be so used.[9]

[8] (21) illustrates a case where the standard used by the adjective is lower than the one determined by the modified noun; (i) illustrates the reverse: B's claim that A's assertion is false is based on a standard for the adjective that is higher than the one correleated with the noun.

(i) A: Kyle's car is an expensive Honda.
B: That's not true! There are no expensive Hondas, only cars that are expensive FOR Hondas.

[9] These observations further show that analyses that attempt to derive the interpretation of predicative uses of gradable adjectives from underlying attributive structures, such as Montague 1974 and Lewis 1970, do not actually have a theoretical advantage over analyses that treat the predicative form as basic. Such approaches crucially assume that a modified nominal always provides a comparison class, so that structures with the form *an A NP* have non-context dependent interpretations. If predicative forms are derived from attributive ones, their context dependence can be explained in terms of principles of ellipsis, not from anything having to do with the semantics of the positive form. That is, on this view the context dependent aspect of *is expensive* involves figuring out whether it is an elided form of *is an expensive*

Turning now to *for*-PPs, these appear to provide a stronger argument for positing a comparison class variable in the semantics. As illustrated by (22), it is contradictory to assert that something is *A for an NP* and simultaneously claim that it has the lowest degree on the *A*-scale for the class of *NPs*, which suggests that the *for*-PP obligatorily determines the standard of comparison.

(22) ??Kyle's car is expensive for a Honda, though it's the least expensive model they make.

Before we draw this conclusion, however, we need to consider a fact that has gone mostly unobserved in discussions of the positive form (though see Wheeler 1972, p. 316): sentences with the structure *x is A for a NP* presuppose that *x* is an NP. For example, all of (23a-c) require Kyle's car to be a Honda, as shown by the infelicity of (24a-c).[10]

(23) a. Kyle's car is expensive for a Honda.
b. Kyle's car is not expensive for a Honda.
c. Is Kyle's car expensive for a Honda?

(24) a. ??Kyle's BMW is expensive for a Honda.
b. ??Kyle's BMW is not expensive for a Honda.
c. ??Is Kyle's BMW expensive for a Honda?

For-PPs contrast in this regard with modified nominals, as illustrated by (25a-c).

(25) a. Kyle's BMW is (really) an expensive Honda.
b. Kyle's BMW is (obviously) not an expensive Honda.
c. Is Kyle's BMW (actually) an expensive Honda?

These examples are somewhat odd because they don't seem to be saying much — (25a) is obviously false under normal circumstances; (25b) obviously true — but they do not involve presupposition failures. The adverbs *really, obviously*, etc. generate contexts that make the exam-

Honda, is an expensive Volkswagen, is an expensive BMW, etc., not calculating a context dependent standard of comparison. If modified nominals do not necessarily provide comparison classes, however, then assuming a derivational relation between the predicative and attributive forms doesn't actually eliminate the problem of computing the standard of comparison in the predicative (or attributive) form.

[10] (ia) is fine, but this can be explained in terms of general principles of presupposition projection, on the assumption that this example contains an implicit *if*-clause, as in (ib).

(i) a. Kyle's BMW would be expensive for a Honda.
b. If Kyle's BMW were a Honda, then it would be expensive for a Honda.

ples informative, but they do not affect the conclusion about the *for*-PPs: inserting them in (24a-c) effects no corresponding improvement in acceptability.

The contrast between *for*-PPs and modified nominals presents a serious challenge to the hypothesis that the comparison class is a semantic argument of the positive form, as implemented in (13). Even if we were to assume, based on (22), that a *for*-PP obligatorily determines the value of the comparison class argument introduced by *pos*, we would still fail to explain the presuppositions of an example like (24a). The interpretation we would assign to this example (taking into account the conclusions of the previous section) would be equivalent to (26), which clearly does not presuppose that Kyle's BMW is a Honda.

(26) Kyle's BMW is at least as expensive as an expensive Honda.

We could augment the analysis by stipluating the presuppositional relation between the comparison class and the argument of the predicate: the latter must be a member of the former. However, this would run into problems with modified nominals. As we saw above, modified nominals are not required to provide the comparison class, but they may provide it. If the argument of the adjective were necessarily presupposed to be a member of the comparison class, then we would predict that an example like (25a) should trigger a presuppositon failure on an interpretation in which Kyle's BMW is asserted to (really) be a Honda whose cost is high relative to other Hondas. In effect, such an interpretation should be blocked, but this is not the case; rather, this is thte most natural interpretation of this example. (27a-b) make the same point, in an even more striking way.

(27) a. ??That mouse is (really) small for an elephant.
b. That mouse is (really) a small elephant.

Since elephants are typically not small, the most natural interpretation of the attributive modifier in (27b) is one in which the standard is computed relative to elephants, which in turn indicates that the nominal is providing the comparison class. But (27b), unlike (27a), does not presuppose that the mouse is an elephant.

To account for these facts, I propose that a *for*-PP has a much more local effect on adjective meanings: it does not provide a comparison class argument for the *pos* morpheme, but rather restricts the domain of the measure function denoted by the adjective to just those objects that are members of the set defined by the nominal complement of *for* (cf. Rusiecki 1985). This is captured by the analysis (28), which treats a *for*-PP as a function of type $\langle\langle e, d\rangle, \langle e, d\rangle\rangle$. I assume here that the

NP part of the *for*-PP contributes a property of individuals, though in principle it could be a set or a kind (see Graff 2000 for discussion).[11]

(28) $[\![[_{\text{PP}} \text{ for a NP}]]\!] = \lambda g \lambda x : [\![\text{NP}]\!](x).g(x)$

Thus if *expensive* is a function from objects (that can have costs) to (positive) degrees of cost, then *expensive for a Honda* is a function from Hondas to degrees of cost:

(29) a. $[\![[_{\text{AP}} \text{ expensive}]]\!] = \lambda x.\mathbf{expensive}(x)$
b. $[\![[_{\text{AP}} \text{ expensive for a Honda}]]\!] = \lambda x : \mathbf{Honda}(x).\mathbf{expensive}(x)$

On this view, the examples in (24) are anomalous because the semantic argument of the adjective (the subject) is not a member of its domain, and composition fails. In contrast, modified nominals do not restrict the domain of the adjectives that modify them, so similar effects do not arise.

If *for*-PPs are modifiers that restrict the domain of a gradable predicate, however, then they cannot also serve as comparison class arguments to *pos*.[12] We earlier concluded that the interpretive effects of modified nominals can be accomodated without positing a comparison class argument, so this leaves only one piece of linguistic evidence for treating comparison classes as semantic arguments of the positive form: the 'bound comparison class' interpretation of examples like (30a), paraphrased in (30b).

(30) a. Everyone in my family is tall.
b. for every x such that x is a member of my family, x has a height greater than the norm for a comparison class based on x.

[11] The property contributed by the NP part is not always merely the denotation of the NP, though it is always a function of the meaning of the NP. For example, the domain restriction contributed by the *for*-PP in (ia-b) is not the property of being a first year graduate student, but rather the property of being work done by a first year graduate student, as illustrated by the anomaly of (ic).

(i) a. This work is quite sophisticated for a first year graduate student.
b. This is sophisticated work for a first year graduate student.
c. ??This second year student's work is quite sophisticated for a first year graduate student.

[12] However, if they are domain restrictors instead of comparison class arguments, this immediately raises the question of why — unlike modified nominals — they must nevertheless be used to compute the standard of comparison. I return to this question below.

In section 2.2, I claimed that the possibility of interpreting (30a) as (30b) indicates the presence of an element in the denotation of the positive form whose value can vary as a function of the denotation of the subject, which I took to be a skolemized comparison class variable (Stanley 2002). However, the analysis of *for*-PPs as domain restrictors suggests an alternative explanation of this phenomenon.

If the domain of a gradable adjective can be explicitly restricted, then it is reasonable to assume that it can be implicitly restricted as well, just like the domains of other functional expressions such as quantificational determiners (von Fintel 1994; Stanley 2000; Stanley and Szabó 2000; Martí 2002; Giannakidou 2004). Of particular relevance is the fact that implicit quantifier domain restrictions can be bound in a manner fully parallel to what we see in (30), as discussed in detail by Stanley 2000.[13] (31a), for example, can have the interpretation in (31b), in which the implicit domain of quantification for *exactly three* is a function of the value of the variable quantified by *most of John's classes*.

(31) a. In most of Johns classes, he fails exactly three students.
b. In most of John's classes x, he fails exactly three students in x.

Let us assume that implicit domain restrictions on gradable adjectives can be bound in the same way. There are a number of different ideas about how implicit domain restriction should be represented at the syntax-semantics interface; for concreteness, I will follow the references cited above and assume that they are representationally encoded. (This assumption is not crucial, however: other implementations of domain restriction would in principle work just as well, as long as they can capture the binding facts.) Specifically, I will assume that English includes a null pronominal *for*-PP with the denotation in (32), where $\mathbf{f}$ is a variable over properties (or sets or kinds; whatever turns out to be most appropriate).

(32) $[\![[_{forPP}\ pro]]\!] = \lambda g \lambda y : \mathbf{f}(y).g(y)$

If $\mathbf{f}$ can be skolemized, then we have an element of the representation other than a comparison class variable whose value could vary as a function of the denotation of the subject: the domain restriction. In other words, a possible interpretation of (30a) is (33), where the first argument of *pos* is an adjective denotation with a domain restriction

[13] Stanley also discusses examples like (30a), but gives an analysis similar to the one I suggest in section 2.2. That is, he assumes that it is a comparison class variable that is bound in (30a), rather than an adjectival domain restriction.

whose value can vary as a function of the values assigned to the variable bound by the quantifier. (Here **momf** is an abbreviation for the property contributed by *member of my family*.)

(33) $\forall x[\mathbf{momf}(x) \rightarrow [\![pos]\!](\lambda y : \mathbf{f}_x(y).\mathbf{tall}(y))(x)]$

The truth conditions associated with (33) are roughly: every member of my family x is tall relative to a domain determined by x. This looks like what we want, but one piece of the puzzle remains: capturing the relation between adjective domain restrictions and the standard of comparison, so that we explain both the fact that the standard in (33) varies with the quantifier, and the fact (discussed above) that a *for*-PP — *qua* explicit domain restriction — obligatorily determines the standard. Capturing this relation leads to a semantic analysis of the positive form that also implements the conclusions of the previous section.

2.5. A VAGUE SEMANTICS FOR THE POSITIVE FORM

The conclusion of section 2.3 was that the point of contextual variation in the positive form cannot merely be a comparison class variable; the function that selects a standard based on a comparison class and an adjective denotation must itself be context dependent, identifying a significant degree of the property expressed by the adjective relative to the interests and attitudes of the participants in the discourse (and potentially other factors as well). The conclusion of section 2.4 was that comparison classes are not actual semantic arguments of the positive form, but rather correspond either to highly salient properties (in the case of modified nominals) or to the explicitly or implicitly restricted domain of the adjective (in the case of *for*-PPs). In fact, there is no reason to assume that comparison classes have any representational status at all: 'comparison class' is merely a descriptive label for whatever property is used to compute the standard of comparison.

Putting these two conclusions together leads to the following hypothesis about the semantics of the positive form: the standard of comparison is determined by a context dependent function from gradable adjectives to degrees, which returns the minimal degree that represents a significant amount of the property measured by the adjective. This degree is determined based both on aspects of the context, such as the interests of the participants in the discourses, discourse salient properties, and so forth, and features of the denotation of the adjective, such as the type of property being measured, and (in particular) the domain of the measure function that the adjective expresses.

This hypothesis is implemented in the semantics of *pos* stated in (34), where **s** is a context dependent function from adjective denotations to degrees with the properties described above: it returns a contextually significant degree of the gradable property measured by the adjective.

(34) $[\![[_{\text{Deg}}\ pos]]\!] = \lambda g \lambda x.g(x) \succeq \mathbf{s}(g)$

The **s** function in (34) corresponds to the DELINEATION FUNCTION of Lewis 1970 and Barker 2002, which maps a gradable adjective meaning to a degree. My proposal expands on the Lewis/Barker analysis in two ways.

First, following Graff 2000 and Bogusławski 1975, I attach more content to **s** than Lewis or Barker attach to their delineation functions. Specifically, the mapping between adjectives and degrees established by **s** is the 'contxtually significant degree' function, which allows me to build Graff's explanation of borderline cases and the Sorites Paradox into the semantics. The former arise from (possibly) incomplete knowledge about the factors that are relevant for determining significance, while the latter stems from the status of **s** as a context sensitive function and from what Graff refers to as the SIMILARITY CONSTRAINT: if two objects are highly similar with respect to some gradable property g, then one has a significant degree of g if and only if the other one does. The result is that whenever two objects that are highly similar with respect to gradable property g are under consideration, the context is modified in such a way that **s** always returns a standard that treats the two objects in the same way (either both fall above the standard or both fall below it). Since such a consideration is involved in evaluating the second premise of a Sorites argument, we are inclined to accept it as true, even if the larger context with respect to which the universal statement is evaluated renders it false.

Further evidence that 'significance' is a matter of semantics — part of the content of the positive form, rather than a by-product of usage or an epistemological illusion — comes from the fact that its effects can be seen even in examples that involve maximally explicit standards of comparison, such as (35a).

(35) a. This novel is long compared to that one (though both are quite short).
b. This novel is longer than that one (though both are quite short).

(35a) can be used to convey the fact that two novels differ in length, even when one or both of them fail to meet the prevailing standard

of comparison for *long*, as shown by the fact that the parenthetical is non-contradictory. The positive form thus has a use that is functionally similar to the comparative form in (35b), an observation that goes back to Sapir 1944. This use can be explained as follows. First, assume that the function of the *compared to* phrase is to ensure that the denotation of the predicate is calculated with respect to a context that includes only the two objects being compared.[14] Second, assume a general requirement on the standard of comparison that it must always support a partitioning of the domain of the predicate into two non-empty sets (a 'positive' and 'negative' extension; see Klein 1980). The result is that if the first novel has a significant degree of length relative to a context that includes just the two novels being compared, it must be longer than the second novel.

There is, however, an important difference between the positive and comparative forms: only the comparative is felicitious in contexts requiring CRISP JUDGMENTS — contexts in which the difference between two objects with respect to the property measured by the adjective are potentially vanishingly small. This is illustrated by the pair in (36).

(36) CONTEXT: A 100 page novel and a 99 page novel.
a. ??This novel is long compared to that one.
b. This novel is longer than that one.

The acceptability of (36b) is unsurprising: the comparative is true of an object if has a degree of the relevant property that exceeds the degree determined compositionally by the linguistic content of the *than*-phrase (see the discussion of the comparative in section 2.1); this relation holds even when the difference between the objects is slight. According to the semantic analysis developed here, the positive form also establishes an ordering between two degrees, but the standard must satisfy the additional requirement of being significant relative to features of the context. In (36a), the context is one in which only the two novels under

[14] Developing an analysis of how exactly this works goes beyond the scope of the paper; however, it is worth observing that that *compared to* differs from a *for*-PP with respect to the presuppositions discussed in the previous section, as illustrated by the contrast in (i).

(i) a. Kyle's BMW is expensive compared to a Honda.
b. ??Kyle's BMW is expensive for a Honda.

If (ib) is bad due to the effect of the *for*-PP on the domain of the adjective, as argued in the previous section, then the fact that (ia) is acceptable suggests that *compared to* is working at a higher level of meaning: instead of modifying the denotation of the adjective itself, it influences the computation of the standard of comparison by imposing constraints on the context in which the positive form is evaluated.

comparison are being considered; since they are extremely similar (with respect to length) the Similarity Constraint ensures that one has a sigificant degree of length if and only if the other one does. But this violates the requirement that the standard of comparison should support a partitioning of the domain of the predicate, resulting in semantic anomaly.[15]

Crucially, if the 'significant degree' component were not a part of the truth conditions of the positive form — if instead, the difference between the positive and comparative forms were merely that the former expressed a ordering between an object and an implicit standard of comparison, while the latter expressed an ordering between an object and an explicit standard of comparison (determined by the *than*-phrase) — this explanation would be unavailable. In a context such as the one in (36), it would be possible to identify the standard of comparison with a degree that supported a partitioning of the domain of the predicate — e.g., the degree of length of the shorter novel — which in turn would derive truth conditions equivalent to those of the comparative form. The fact that this is impossible demonstrates that there is a semantic difference between the comparative and positive forms that goes beyond the explicit/implicit standard distinction. On the analysis proposed here, this is the requirement that the implicit standard of comparison used by the positive form must be a significant degree of the relevant property.

The second difference between my approach and the delineation function approach of Lewis and Barker is that I assume a more prominent role for the adjective denotation in the computation of the standard. In the analyses developed by Lewis and Barker, the role of the adjective denotation in the computation of the standard is merely to provide information about the relevant scale (i.e., a standard for *tall* should be a degree on the height scale, while one for *expensive* should be a degree on the cost scale). However, as we saw in section 2.4, the relation between adjective denotation and standard is much stronger, since changing one feature of adjective meaning — the domain of the function it expresses — has an impact on the standard of comparison.[16]

[15] Alternatively, we could trace the anomaly of (36a) to the fact that if the speaker is obeying the domain partitioning requirement, then in uttering (36a) he commits himself to the unlikely position that a length of 99 pages is not significant (for a novel in the context of utterance) while a length of 100 pages is. Thanks to Ede Zimmermann for pointing out this interpretation of the data to me.

[16] A third difference is that both the Lewis and Barker analyses treat gradable adjectives as expressions of type $\langle e, t\rangle$: neither assumes the type 'degree' (though Barker does definie the truth conditions of gradable adjectives in terms of degrees). A consequence of this assumption is that the semantics of comparatives and other types of degree modifiers must be characterized in terms of (modifications to) the

One illustration of this relation is the binding facts, which can now be straightforwardly explained without a comparison class variable. Assuming as above that an adjectival domain can be implicitly restricted via a skolemized domain restriction variable, (37b) is a possible interpretation of (37a).

(37) a. Every member of my family is tall.
b. $\forall x[\mathbf{momf}(x) \rightarrow [\lambda y : \mathbf{f}_x(y).\mathbf{tall}(y)](x) \succeq \mathbf{s}(\lambda y : \mathbf{f}_x(y).\mathbf{tall}(y))]$

What varies in (37b) is not a comparison class, but rather the domain of the function expressed by the adjective, which is restricted to things that satisfy the contextual skolem function $\mathbf{f}_x$ for each value of x determined by the subject. In a particular context, the value of this domain restriction might be something like 'people that are the same age as x' (which is of course going to be satisfied by any value of x). Crucially, changing the domain means changing the function expressed by the adjective. Thefore, since **s** computes a standard of comparison based on this function (as well as the various contextual factors relevant to establishing significance), changes in the domain based on the value of x, allow for corresponding changes in the standard of comparison based on x.[17]

A more striking illustration of the relation between adjective denotation and the standard of comparison is the fact that a restricted

---

delineation function, a feature that is shared by other analyses that eschew degrees (e.g., McConnell-Ginet 1973; Kamp 1975; Fine 1975; Klein 1980). (Not all of these analyses posit a delineation function, but they include the functinonal equivalent.) This is not the case in degree-based analyses, in which it is possible to define the interpretation of gradable predicates strictly in terms of relations between degrees (as illustrated in section 2.1 for simple phrasal comparatives), independent of the delineation function. This difference is irrelevant to the issues under consideration here (since the two approaches agree that the semantics of the positive form is stated in terms of a contextual mapping from adjectives to degrees), but may lead to different predictions about the semantics of comparatives based on data like (36) (Kennedy 2005).

Finally, Barker's (2002) analysis is explicitly dynamic, though the dynamic component could be overlaid on the analysis developed here with no significant change to the core proposals.

[17] This account of variability in the standard of comparison based on quantifiers carries over directly to examples involving 'sloppy identity' of the standard in ellipsis constructions such as (i) (Klein 1980; Ludlow 1989).

(i) That elephant is large, and that flea is too.

(i) has a reading in which the elephant is claimed to be large for an elephant (or more accurately, for 'things like it') and the flea is large for a flea. Under the assumptions outlined here, this is just another case of binding into an implicit adjectival domain restriction, with consequent affect on the value of the standard of comparison.

adjectival domain must be used to compute the standard of comparison, as discussed in section 2.4. This is illustrated by the following example.[18] Consider a context in which an experiment is being run involving a subject who is watching creatures moving around a maze and describes what is happening. There is a mix of different kinds of animals — insects, reptiles, and mammals — but nothing is bigger than a small mouse. Sometimes a creature leaves the maze, sometimes a new one enters. Suddenly, a rat-like creature of substantial size (especially compared to the creatures currently in the maze) enters. Assume that the subject first says (38).

(38) Another animal just entered the maze.

Now consider the following possible continuations:

(39) a. It's a mammal, and it's large. It might be a rat.
b. It's a large mammal. It might be a rat.
c. ??It's large for a mammal. It might be a rat.

(39a) is clearly felicitous in this context, indicating that the standard of comparison for non-explicitly restricted *large* is a function of the sizes of the objects in the experiment. (39b) is also possible, further demonstrating that modified nominals do not have to provide the property on the basis of which the standard of comparison is computed. (39c), however, is infelicitous because it is obviously false: even a rat (or a rat-like creature) of subtantial size is not large for a mammal. Thus even in a context in which other properties are highly salient (such as the property of being an object in the maze), a restricted adjectival domain must be used to calculate the standard of comparison.

The semantic analysis of the positive form in (34), which claims that the standard is computed as a function of adjective denotation, provides a basis for capturing this correlation. According the the analysis of *for*-PPs developed in section 2.4, the crucial difference between (39a-b)

[18] The example that I used to illustrate this relation in section 2.4 was (i), which shows that it is contradictory to simulaneously assert that an object is *A for an NP* and that it has the lowest value on the *A*-scale for the set of NPs.

(i) ??Kyle's car is expensive for a Honda, though it's the least expensive model they make.

This particular example can actuallly be explained in terms of general constraints on informativity: if the least expensive element of the domain of the predicate counts as having the problem in question, then the sentence could not fail to be true. In order to ensure that a particular assertion is informative, then, it must be the case that the standard of comaparison is such that the predictate is false of at least some of the objects in the domain (cf. Klein 1980).

and (39c) is that only in the latter is the domain of the adjective *required* to be the set of mammals, via the semantics of the explicit domain restrictor (the *for*-PP). In contrast, although (39a-b) could in principle have implicit restrictions with equivalent semantic effects, the consequence of such a restriction is obviously false truth conditions. As a result, these interpretations are ignored in favor of ones in which the input to the standard-identifying function **s** is the unrestricted adjective denotation, and significance is calculated in terms of other contextual factors. Since **s** is a function from adjective denotations to degrees, and restricting an adjective's domain entails a change in its denotation (a function from objects in general to degrees of largeness is distinct from a function from mammals to degrees of largeness), we expect explicit domain restriction to influence the computation of the standard of comparison, exactly as we see in (39c) vs. (39a-b).

However, even though the semantic analysis in (34) ensures that, all other things being equal, a change in the domain of an adjective will influence the computation of the standard of comparison, it does not *require* that once the domain is fixed the standard will end up being a significant degree of the relevant property *based only on the domain of the adjective*. That is, even if a restricted adjective denotation is obligatorily part of the calculation of the standard, as the analysis predicts, the inherent context dependence of **s** doesn't exclude the possibility that other salient contextual information could also come into play, especially if the domain restriction alone resulted in an infelicitous interpretation. If this were correct, then it ought to be possible to understand (39c), for example, as an assertion that the animal entering the maze has a size that is significant for a rat-like mammal. The fact that (39c) cannot mean this clearly indicates the action of some constraint prohibiting anything but the restricted domain from being used to compute the standard.

For the moment, I will simply state the constraint descriptively as in (40).

(40) *Domain Dependence*
If a gradable adjective has a restricted domain, then the standard of comparison must correspond to a significant degree of the relevant property relative to the domain.

One plausible explanation for this constraint is a functional one: why bother to restrict the domain if not to specify the basis on which the standard of comparison should be computed? Although I suspect that the ultimate source of this constraint involves functional considerations, I will argue in section 4 that the correlation between adjectival domain

and standard of comparison is one instance of a more general principle of 'interpretive economy', which in the case of the positive form of a gradable predicate requires the truth conditions to be computed strictly on the basis of features of the conventional meaning of the adjective if possible. Such features include the domain of the function expressed by the adjective, deriving the correlation observed here.

To get a full picture of the empirical evidence for such a principle, however, we must turn to the case of absolute gradable adjectives, which on the surface appear to raise significant problems for the semantic analysis of the positive form that we have arrived at with here.

3. Absolute gradable adjectives

3.1. The problem

Most of the literature on vagueness assumes (implicitly or explicitly) that all gradable predicates in the positive form have the properties that we have observed and analyzed for adjectives like *expensive*, *tall*, *large*, etc.: truth conditional variability, borderline cases, and Sorities sensitivity. And indeed, if the analysis of the positive form developed in section 2.5 is correct, then any gradable adjective — any expression of type $\langle e, d\rangle$ — is predicted to show these characteristics in the positive form, since they follow from the semantics of the positive degree morpheme *pos* repeated in (41).

(41) $\quad [\![[_{\text{Deg}}\ pos]]\!] = \lambda g \lambda x.g(x) \succeq \mathbf{s}(g)$

Specifically, they follow from the assumption that the standard function $\mathbf{s}$ returns a degree on the g-scale that represents a significant amount of the property measured by g relative to salient features of the context (which include the comparison class, though that term should now be understood in a purely descriptive sense as the label for whatever property is relevant for establishing significance — either the domain of g or some other contextually salient property).

In fact, this prediction is incorrect. In addition to the large class of gradable adjectives that show these features of vagueness — henceforth RELATIVE gradable adjectives — there is a well-defined set of adjectives that are demonstrably gradable but do not have context dependent interpretations, do not give rise to borderline cases, and do not trigger the Sorites Paradox — at least not in the way we have seen so far.

Following Unger 1975 and more recently Kennedy and McNally 2005, I will refer to this class as ABSOLUTE (gradable) adjectives.[19]

Absolute adjectives come in two varieties. MINIMUM STANDARD absolute adjectives, such as those in (42), simply require their arguments to possess some minimal degree of the property they describe; they do not require the degree to which the arguments possess this property to be significant based on a comparison class.

(42) a. The gold is impure.
b. The table is wet.
c. The door is open.
d. The rod is bent.

Under normal usage, (42a) does not mean that the degree to which the the gold is impure exceeds some contextual standard of impurity (for gold); it simply means that the gold contains some amount of impurity. Likewise, (42b) is true as long as there is some amount of water on the table; (42c) just requires some minimal positive aperture of the door; and (42d) is true of a rod that has a non-zero degree of bend.

MAXIMUM STANDARD absolute adjectives such as those in (43) require their arguments to posses a maximal degree of the property in question.

(43) a. The platinum is pure.
b. The floor is dry.
c. The door is closed.
d. The rod is straight.

(43a) typically means that the platinum is totally pure, not that its contents fall above some context dependent standard of 'significant purity'; (43b) is an assertion that the floor has no moisture on it; (43c) requires the door to be completely closed; and (43d) requires a completely straight rod.

[19] Predicates formed out of absolute adjectives have also been studied in detail by Rotstein and Winter 2004, who refer to adjectives like those in (42) 'partial' predicates and adjectives like those in (43) 'total' predicates (using terminology introduced in Yoon 1996). Rotstein and Winter develop semantic analyses of these predicates in terms of their underlying scalar properties, a strategy that I will also endorse in section 4. They do not address the relation between relative and absolute adjectives, however, or the question of whether it is possible to develop a fully general analysis of the positive form which assigns the correct truth conditions (relative, absolute miniumum, absolute maximum) given particular adjectives as inputs, which is what I am trying to do here.

Clear evidence that absolute gradable adjectives are gradable comes from the fact that they are perfectly acceptable in comparatives and with other degree morphology, as shown by (44)-(45).

(44) a. The platinum is less impure than the gold.
b. The table is wetter than the floor.
c. The door isn't as open as I want it to be.
d. This rod is too bent to be of use for this purpose.

(45) a. The gold is less pure than the platinum.
b. The floor is dryer than the table.
c. The door is closed enough to keep out the light.
d. This rod is too straight to be of use for this purpose.

They contrast in this regard with true non-gradable adjectives, which are anomalous in comparatives:[20]

(46) a. ??The platinum is less geological than the gold.
b. ??The table is more wooden than the floor.
c. ??The door isn't as locked as I want it to be.
d. ??This rod is too hand-made to be of use for this purpose.

The acceptability and interpretation of absolute adjectives in comparatives and other degree constructions indicates not only that they have the same semantic type as relative gradable adjectives — both denote functions of $\langle e, d\rangle$, and so can combine with degree morphology — but also that they have fundamentally the same kinds of meanings: they take an object and return a measure of the degree to which it possesses some gradable property. For example, in order to accurately capture the meaning of the comparative predicate in (45b), we simply need to assume that *dry* denotes a function **dry** from objects to degrees of dryness (or perhaps more accurately, 'absence of moisture content'). As illustrated in (47) (where the interpretation of the comparative morpheme *more* is the same as in the example discussed in section 2.1), this derives a meaning for the comparative that is exactly what we want: it is true of an object if it has a degree of dryness that exceeds that of the table.

(47) $[\![\text{more}]\!]([\![\text{dry}]\!])([\![\text{than the table}]\!])$
$= [\lambda g \lambda y \lambda x.g(x) \succ g(y)](\lambda z.\mathbf{dry}(z))(\mathbf{the\ table})$

[20] It is often possible to coerce a gradable interpretation from a non-gradable adjective (e.g., (46b) might be understood to mean that the table has more wood in it than the door), but such interpretations are clearly marked. In contrast, the comparatives in (44)-(45) are perfectly natural and felicitous.

$= [\lambda y \lambda x.\mathbf{dry}(x) \succ \mathbf{dry}(y)](\mathbf{the\ table})$
$= \lambda x.\mathbf{dry}(x) \succ \mathbf{dry}(\mathbf{the\ table})$

However, if absolute adjectives have the same semantic type and the same kind of meaning as relative adjectives, their interpretations in the positive form are unexpected. Since they are expressions of tyep $\langle e, d\rangle$, they must combine with degree morphology to derive a property of individuals. According to the analysis of the positive form developed in 2, unmarked gradable adjectives combine with *pos*, which has the denotation repeated above in (41). This means that the predicates in (42a) and (43a), for example, should have the compositional analyses and truth conditions shown in (48a-b).

(48) a. $[\![pos]\!]([\![\text{impure}]\!])$
$= [\lambda g \lambda x.g(x) \succeq \mathbf{s}(g)](\lambda y.\mathbf{impure}(y))$
$= \lambda x.\mathbf{impure}(x) \succeq \mathbf{s}(\lambda y.\mathbf{impure}(y))$

b. $[\![pos]\!]([\![\text{pure}]\!])$
$= [\lambda g \lambda x.g(x) \succeq \mathbf{s}(g)](\lambda y.\mathbf{pure}(y))$
$= \lambda x.\mathbf{pure}(x) \succeq \mathbf{s}(\lambda y.\mathbf{pure}(y))$

(48a) is true of an object if it has a contextually significant degree of impurity, and (48b) is true of an object if it has a contextually significant degree of purity. The problem is that these denotations do not accurately capture the truth conditions described above for (42a) and (43a): (48a) is too strong, and (48b) is too weak. More generally, any analysis that is designed to provide a general, comprehensive analysis of the properties of vagueness exhibhited by the positive form of relative gradable adjectives will fail to provide an accurate account of the interpretation of absolute gradable adjectives in the positive form. In short: if vagueness arises from the semantics of the positive form (whether from the meaning of a *pos* morpheme, as in the analysis proposed in section 2, or from a rule of default interpretation), then absolute adjectives should be just as vague as their relative counterparts. As the next sections will demonstrate in detail, this prediction is incorrect.[21]

[21] At this point it might appear that an analysis in which vagueness is a feature of adjective meaning generally rather than the positive form specifically (as in Lewis 1970, Kamp 1975; Fine 1975; Klein 1980; Barker 2002) would be better equipped to handle the data, since it could be captured as a lexical distinction. This conclusion would be incorrect, however. Such an account would involve analyzing absolute adjectives as fixed functions from individuals to truth values, which would in turn predict that they should be unacceptable in comparatives, since the semantics of comparison in such accounts involves quantifying over the set of possible interpretations of a vague predicate. In other words, in such accounts gradability entails vagueness, just as in the analysis of the positive form developed in section 2. Whether such analyses can take advantage of the proposals I will make in section 4 to derive

3.2. Evidence for the absolute/relative distinction

3.2.1. *Imprecision vs. vagueness*

The problem of absolute adjectives has not been explicitly addressed in previous work on vagueness or the semantics of gradable predicates, possibly because there is a strong initial intuition that the adjectives in (42) actually require something significantly more than a minimum standard, and that those in (43) actually allow something less than a maximum standard. These intuitions are supported by examples like those in (49).

(49) a. I'm not awake yet.
b. The theater is empty tonight.

Awake is a minimum standard adjective, but (49a) can be felicitously uttered by someone who is not talking in his sleep. Similarly, *empty* is a maximum standard adjective, but (49c) can be used to describe a situation in which only a very few people show up to a film in a very large movie theater. These examples appear to call into question the empirical claims made in the previous section.

In fact, however, these examples illustrate a phenomenon that is distinct from vagueness, though typically exists alongside it: IMPRECISION. As discussed by Pinkal (1995), there are many expressions that have imprecise uses, but which are not vague. One clear example is predicates formed out of relative gradable adjectives and measure phrases, such as (50).

(50) The rod is 10 meters long.

(50) can be felicitously used to describe a rod whose actual length falls somewhere close to 10 meters, in a range that is itself be subject to contextual variation (e.g., 995cm to 1005cm in one context; 9998mm to 10,002mm in another, and so forth).

The difference between *10 meters long* and *long* is that the former, but not the latter, allows for what Pinkal (1995, pp. 99-100) refers to as NATURAL PRECISIFICATIONS: it is possible to construct a context in which *10 meters long* cleanly distinguishes between objects based on potentially very slight differences in length (e.g., a scientific experiment or a construction project); it is difficult (if not impossible) to do the same for the simple positive *long*. Put another way, it is possible to construct a natural context for the former but not the latter in which

---

the relative/absolute distinction depends on whether or not scales and degrees play a role in the semantics of gradable predicates (as in e.g. Barker 2002); see section 4.4 for discussion.

borderline cases are eliminated. This is illustrated by the examples in (51).

(51) a. We need a 10 meter long rod for the antenna, but this one is 1 millimeter short of 10 meters, so unfortunately it won't work.

b. ??We need a long rod for the antenna, but since *long* means 'greater than 10 meters' and this one is 1 millimeter short of 10 meters, unfortunately it won't work.

(51a) is perfectly natural, especially if we know the context is one in which precision is important, and small differences of measurement can make a large difference in outcome (such as building a spacecraft). (51b), however, is decidedly unnatural, even in the same type of context. It is not impossible to make sense of this example, but as Pinkal points out, forcing *long* to be interpreted in this way involves fundamentally changing the meaning assigned to the positive form by the semantics.

Absolute adjectives behave like measure phrases with respect to natural precisifications. For example, assume that I am a detective in search of a violent criminal, and I'm trying to find out whether he might be hiding out in a particular movie theater. In the context described above (very few people watching a popular movie), I would consider the projectionist to be lying if he uses (49b) to respond to my question *Is anyone in the theater tonight?*. Likewise, (52) is perfectly natural in the spacecraft construction context, even though it implies that a small amount of bend is enough to prevent the rod from counting as straight.

(52) The rod for the antenna needs to be straight, but this one has a 1 mm bend in the middle, so unfortunately it won't work.

I conclude from these observations that imprecise uses of absolute adjectives in the positive form do not call into question the central descriptive claim that such expressions have truth conditions that make reference to fixed (maximal or minimal) standards of comparison. On the contrary, they provide evidence in support of this conclusion, by highlighting the contrast between relative and absolute adjectives with respect to the possibility of natural precisifications. In the following sections, I will outline four additional sets of facts that provide support for the distinction between relative and absolute adjectives in the positive form and for the specific claim that the crucial difference between them has to do with the standard of comparison: whether it is contextually

variable or whether it is (or at least defaults to) an endpoint on a scale.[22]

3.2.2. *Entailments*

Clear evidence for the relative/abolute distinction comes from entailment patterns. (The facts in this section are also dicussed in Kennedy and McNally 2005.) If the standards of comparison associated with the positive forms of absolute adjectives are minimal or maximal degrees (depending on the adjective), the truth conditions of the positive can be stated as in (53a) for a minimum standard adjective g_{min}, and (53b) for a maximum standard adjective g_{max}.[23]

(53) a. $\lambda x.g_{min}(x) \succ min(\text{SCALE}(g_{min}))$
b. $\lambda x.g_{max}(x) = max(\text{SCALE}(g_{max}))$

For example, plugging the measure functions expressed by *impure* and *pure* in for g_{min} and g_{max}, respectively, derives properties with precisely the truth conditions that I described for the predicates in (42a) and (43a) in section 3.1:

(54) a. $\lambda x.\mathbf{impure}(x) \succ min(\text{SCALE}(\lambda y.\mathbf{impure}(y)))$
b. $\lambda x.\mathbf{pure}(x) = max(\text{SCALE}(\lambda y.\mathbf{pure}(y)))$

These interpretations make specific predictions about entailments of the positive form of absolute adjectives that are distinct from those made for the positive form of relative adjectives. First, (54a) predicts that a negative assertion *x is not* A_{min} should entail that *x* possesses

[22] A full explanation of imprecision goes well beyond the scope of this paper (though I return to it briefly in section 4.4), but a couple of potential approaches immediately suggest themeselves. One possibilty would be to explain imprecision in terms of more general pragmatic principles governing the interpretation of 'loose talk'; see Kennedy and McNally 2005 for a suggestion along these lines that makes use of Lasersohn's (1999) theory of PRAGMATIC HALOS. A related possibility, suggested in Pinkal 1995, would be to assume that uncertainty in measurement systems is reflected in the interpretation of otherwise precise expressions (like measure phrases and absolute adjectives) by a contex-dependent notion of 'tolerance' of application of a predicate. Finally, we could build uncertainty in measurement directly into the scalar representations and keep the semantics precise by by analyzing precise vs. imprecise uses of gradable predicates in terms of different granularities of degrees, so that e.g. the same container could count as maximally empty at coarse granularities but not at finer ones. Whatever explanation we adopt, it should extend to an account of imprecise uses of absolute predicates, such as those illustrated in (49).

[23] Recall that I am assuming that gradable adjectives denote functions from objects to degrees, and that an ordered set of degrees is a scale. The SCALE function in (53) is therefore just a function from a gradable adjective meaning to its range. I dicuss this issue in more detail in section 4.

no amount *adj*-ness at all, assuming that the minimal degree on a scale represents a zero amount of the relevant property. The contradictory statements in (55) illustrate that this prediction is borne out. (# is used to indicate contradiction.)

(55) a. #The gold is not impure, but there is some lead in it.
b. #My hands are not wet, but there is some moisture on them.
c. #The door isn't open, but it is ajar.

Second, (55b) predicts that an assertion of *x is* A_{max} should entail that x has a maximal amount of 'A-ness', i.e., that x cannot be more A than it is. This sort of entailment is difficult to test, since maximum standard adjectives readily allow imprecise uses. However, as observed by Unger (1975), it is possible to force a precise interpretation by adding focal stress (specifically a falling tone) to the adjective. When we do this, as in (56), we see that the expected entailments arise:

(56) a. #My glass is FULL, but it could be fuller.
b. #The line is STRAIGHT, but you can make it straighter.

In contrast to absolute adjectives, the truth conditions for a relative adjective in the positive form require that its argument falls above whatever degree represents a significant amount of the property for the context. As a result, neither of the above entailments should hold: negation should be compatible with a positive (if less than significant) degree of the measured property, and assertion should not rule out higher values. This is correct:

(57) a. Sam is not tall, but his height is normal for his age. (requires Sam to have some degree of tallness)
b. That film is interesting, but it could be more interesting.

A related argument involving entailments is discussed in Cruse 1986 (see also Rotstein and Winter 2004). As shown by the examples in (58), there exist pairs of antonyms such that negation of one form entails the assertion of the other:

(58) a. The door is not open. ⇒ The door is closed.
b. The table is not wet. ⇒ The table is dry.
c. The baby is not awake. ⇒ The baby is asleep.

The explanation for this is straightforward: both members of the pairs in (58) are absolute adjectives, but the positive adjectives impose minimum standards while the negative adjectives impose maximum standards. Since a minimal positive degree corresponds to a maximal neg-

ative degree on the same scale (see Kennedy 2001), the entailment relations in (58) follow from the truth conditions in (53).

Relative antonyms do not show the same entailment relations:

(59) a. The door is not large. $\not\Rightarrow$ The door is small.
b. The table is not expensive. $\not\Rightarrow$ The table is inexpensive.
c. The baby is not energetic. $\not\Rightarrow$ The baby is lethargic.

Again, this follows from the fact that the standards for both positive and negative relative gradable adjectives are contextually identified and based on adjective denotation. The fact that e.g. *large* and *small* both measure size ensures that in any context, there should be some relation between their respective standards (in particular, the standard of largeness for an object should never be less than the standard of shortness for the same object), but they need not be the same degree. This allows for the possibility of a 'grey area' between the standards onto which fall objects that are neither large nor small (Sapir's (1944) 'zone of indifference'; Klein's (1980) 'extension gap').

We also see differences in the entailments of relative and absolute adjectives in comparative constructions. Assuming that the comparative imposes an asymmetric ordering on its arguments (see the truth conditions for the comparative in (5b)) and that the truth conditions associated with minimum and maximum standard absolute adjectives in the positive form are as in (53), we predict that comparatives with absolute adjectives should generate positive and negative entailments to the positive form, respectively, depending on whether we have a minimum or maximum standard adjective (cf. Rusiecki 1985). This prediction is borne out, as shown by the examples in (60)-(61).

(60) a. The floor is wetter than the countertop. $\Rightarrow$
b. The floor is wet.

(61) a. The floor is drier than the countertop. $\Rightarrow$
b. The countertop is not dry.

(60a) is true only if the floor has some degree of wetness: if it had zero wetness, then it could not possibly have a greater degree of wetness than the countertop.[24] This satisfies (53a) and generates the entailment to

[24] I assume that (60a) does not entail that the countertop also has some degree of wetness — it could in principle be completely dry (zero wetness). There is however a strong implicature that the countertop is also wet. This can be explained as follows: if the countertop had zero wetness, then the truth conditions of (60a) would be identical to the positive statement in (60b). The fact that the more complex comparative form is used — and in particular, the fact that the countertop is explicitly introduced as a reference point for characterizing the wetness of the floor — implicates that

(60b). Similarly, in order for (61a) to be true, it must be the case that the countertop is not maximally dry (though the floor might be). If the standard for dryness is the maximum value on the scale, as stated in (53b), then it follows that the countertop is not dry.

In comparison, a canonical property of comparatives with relative adjectives, is that they do not give rise to positive or negative entailments in the comparative form, as illustrated by (62)-(63).

(62) a. Rod A is longer than rod B. $\not\Rightarrow$
b. Rod A/B is (not) long.

(63) a. Rod A is short than rod B. $\not\Rightarrow$
b. Rod A/B is (not) short.

This follows from the semantics of comparison and the semantics of the positive developed in section 2.5: the mere fact that one object exceeds another with respect to some relative property tells us nothing about how the objects stand in relation to a contextually significant amount of the relevant property.

3.2.3. *For-PPs*

A second, somewhat more variable, difference between relative and absolute adjectives involves the acceptability of *for*-PPs. We have already seen that such expressions are used to fix the domain of a gradable adjective, thereby affecting the computation of the standard of comparison in accord with the Domain Dependence principle stated in (40).

(64) a. The baby is {tall, short, fast, talkative} for a two year old.
b. That table is {small, sturdy, unusual} for a piece of outdoor furniture.
c. That glass is {expensive, clean, dirty} for a wine glass.
d. The door is {strong, big, wide} for an office door.

For-PPs are often infelicitous with absolute adjectives, however, as illustrated by the examples in (65).

(65) a. ??The baby is {awake, asleep} for a kid who hasn't napped all morning.
b. ??That table is {wet, dry} for a piece of outdoor furniture.

(60a) provides more information than the simpler positive form. But this will only be the case if the countertop has some (observable) degree of wetness itself. Note that this inference is not a presupposition: as shown by (i), it disappears under negation.

(i) The floor is not wetter than the countertop; they're both perfectly dry.

c. ??My glass is {full, empty} for a wine glass.
d. ??The door is {closed, open} for an office door.

This difference follows if the interpretation of the positive form of an absolute adjective involves a fixed (maximal or minimal) standard of comparison: restricting the domain of the adjective adds nothing to the interpretation of the predicate, so the *for*-PPs serve no semantic function.

This restriction is not absolute, however. The examples in (66) involve absolute adjectives, but the *for*-PPs are felicitous. (I am grateful to Jeff King for bringing these to my attention.)

(66) a. That cue is straight for a pool cue in a dive like this.
b. This theater is empty for a theater showing a popular movie.

These facts show that it is in fact possible to use an explicit domain restriction to shift the standard of a maximum standard absolute adjective, in effect deriving a relative interpretation. (It is more difficult to construct similar examples for minimum standard adjectives.) However, I do not think that these examples call into question the relative/absolute distinction; rather, they can be used to reinforce the initial claim that such a distinction should be made.

Observe that (66a-b) strongly implicate the negations of the bare positives in (67). (66b), for example, would be infelicitous if used to describe a theater that is in fact completely empty.

(67) a. That cue is not straight.
b. This theater is not empty.

This follows if the standards for the absolute adjectives in these examples at least default to maximum values. If the speaker wants to convey that a particular pool cue is straight, he can simply use the bare positive form. The fact that a *for*-PP is used in (66a) to provide an explicit domain therefore implicates that the truth conditions associated with the bare positive do not obtain, leading to the inference in (66a).

In contrast, relative adjectives with *for*-PPs do not generally give rise to negative inferences to the positive form. Neither of the sentences in (68) implicate their counterparts in (69), for example.

(68) a. Lyosha is old for a dog.
b. Sterling is large for a 9 month old.

(69) a. Lyosha is not old.
b. Sterling is not large.

This follows from our assumption that there is no default standard associated with the positive form of a relative adjective; that it is instead context dependent. Lyosha could be both old for a dog and old with respect to whatever property is relevant to the evaluation of (69a) in the context of utterance (which may in fact be the property of being a dog).

3.2.4. *Standards of differentiation*

A striking example of the context sensitivity of relative adjectives — in particular, of the way that the standard of comparison can be shifted based on salient contextual information — comes from experimental work demonstrating that a relative adjective may be used as part of a definite description to distinguish between two objects that differ only in the extent to which they possess the property expressed by the adjective, even when neither or both objects would be judged to have the property (to a significant degree) outside of the differentiation task (Sedivy et al. 1999; Syrett et al. 2005; see aslo Kyburg and Morreau 2000).

For example, Syrett et al. 2005 describe an experiment comparing adults' and children's ability to contextually shift a standard of comparison in which subjects are presented with examples like (70) as requests for one of two objects, both of which have been judged to be either not long or long in an independent task.

(70) Please give me the long one.

Both children and adults systematically accept these examples as requests for the longer of the two objects, which is fully expected if the positive form of a relative adjective involves a relation to a context dependent standard of comparison. In particular, two contextual factors are at work here: the high salience of the two objects under consideration and the uniqueness and existence presuppositions of the definite description. In order to satisfy the latter, a standard should be computed that makes the description true one object but false of the other. This can be done by taking advantage of the former, and identifying the standard as some degree that maximally differentiates between the two objects, e.g., the mean of the distance between them.

What is important to the current discussion is that (in adult grammar, at least) absolute adjectives do not permit this sort of use. (71), which contains the maximum standard absolute adjective *full*, was accepted in a context involving a full jar and a partially full jar, but it was rejected (uniformly by adults; partially by children) in a context involving two partially full jars, where one was clearly fuller than the other.

(71) Please give me the full one.

(72) Please give me the spotted one.

(73), with the minimum standard adjective *spotted*, was accepted in a context involving a spotted disc and a spotless disc, but was rejected (uniformly by both adults and children) in a context involving a disc with just a few spots and one with many spots.

These results strongly support the claim that absolute adjectives have endpoint-oriented, fixed standards of comparison. The problem with (71) in the 'non-full/less full' context is that the existence presupposition of the definite is not met: neither jar is full. The problem with (72) in the 'spotted/more spotted' context is that the uniqueness presupposition of the definite is not met: both discs are spotted. Both of these problems could be resolved, and the sentences made felicitous, if the standards of comparison for the respective adjectives could be shifted in the same way we saw for *long* above. The fact that the sentences remain infelicitous, despite the context, shows that the standards for *full* and *spotted* are maximum and minimum values on the scale, respectively.[25]

3.2.5. *The Sorites Paradox*

A final difference between relative and absolute adjectives involves their interaction with the Sorites Paradox. As we have seen, a defining char-

[25] The same set of experiments demonstrated that both children and adults accept comparative forms of absolute and relative adjectives in the task regardless of whether the truth conditions for the positive are met: all of the examples in (i) are treated as felicitous in each of the conditions mentioned above (two long rods; two short rods; full/non-full container; non-full/less full container; spotted/more spottted disc; non-spotted/spotted disc).

(i) a. Please give me the longer one.
b. Please give me the fuller one.
c. Please give me the more spotted one.

For adults, this is expected: the semantics of the comparative involves a relation to an arbitrary degree (in this case, the degree to which the less long/full/spotted object possesses the property), not a relative/absolute standard of comparison. For children, this was interesting because it demonstrated that they are interpreting the positive form of relative adjectives in an adult-like way. That is, they are shifting a contextual standard of comparison, rather than coercing a comparative meaning out of the positive form. If they were doing the latter, then their behavior on *full* and *spotted* would have mirrored their behavior on *fuller* and *more spotted*.

What remains to be explained is the fact that showed a significantly high tendency to interpret maximum standard adjectives relatively, but they were fully adult-like on minimum standard adjectives.

acteristic of relative adjectives is that they give rise to the Paradox. This is illustrated again for the relative adjective *big* in (73).

(73) P1. A theater with 1000 seats is big.
P2. Any theater with 1 fewer seat than a big theater is big.
C. Therefore, any theater with 10 seats is big.

Building on Graff 2000, I suggested in section 2.5 that the paradox derives from semantics of the positive form, in particular from the properties of the contextual parameter of evaluation **s**, which maps an adjective denotation to a contextually significant standard of comparison. The reason that we fail to reject the second premise of the Paradox is that whenever we consider two objects that are highly similar with respect to gradable property g to determine if one has a significant degree of the property but the other doesn't, we modify the context in such a way that **s** always returns a standard that treats the two objects in the same way.

If the positive form of an absolute adjective also expressed a relation to a contextually significant degree, then it ought to also give rise to the Sorites Paradox. In particular, we should show the same willingness to accept the second premise of the paradox that we show with relative adjectives like *big*. This is not the case, however: (74), for example, does not lead to a paradoxical conclusion, precisely because the second premise is quite naturally judged to be false.

(74) P1. A theater in which every seat is occupied is full.
P2. Any theater with one fewer occupied seat than a full theater is full.
C. Therefore, any theater in which half of (none of, etc.) the seats are occupied is full.

Though we might be willing to accept the second premise in certain contexts (where precision is not important; see the discussion in section 3.2.1), it is just as easy to construct a context in which we would reject it.

(75) makes the same case for the minimum standard adjective *impure*.

(75) P1. Water that contains some amount of contaminants is impure.
P2. Water that contains fewer contaminants than impure water is impure.
C. Water that contains no contaminants is impure.

The first premise is clearly true, but the second is again easily judged false: one way for a quantity of water to have fewer contaminants than impure water would be to have no contaminants at all, but in that case the water would no longer be impure!

3.3. Summary

The conclusion to be drawn from the facts discussed in this section is that at the very least, we need to assume that the standard of comparison in predicates headed by absolute adjectives 'defaults' to a minimum or maximum value: it need not be a contextually determined significant degree of the relevant property, and in the vast majority of cases, it clearly isn't. Although it is possible to force a relative standard in some contexts, this is the marked case. The bulk of the evidence — the possibility of natural precisifications, entailment patterns, the rigidity of the standard, the failure to trigger the Sorites Paradox — points to the conclusion that absolute adjectives in the positive form use fixed, maximal or minimal (depending on the predicate) standards of comparison.

But this in turn means that a semantics of the positive form that is stated in terms of an ordering with respect to a 'contextually significant degree' of the relevant property, such as the one designed to account for the facts in section 2, can work only for relative adjectives. Such an analysis both derives the wrong truth conditions for absolute adjectives and fails to explain the differences between relative and absolute adjectives documented here.

4. Vagueness and grammar

4.1. The polysemy of the positive form

The descriptive conclusion to draw from the facts discussed in the preceding sections is that the positive form can have (at least) the three distinct interpretations specified in (76), depending on the kind of adjective that heads the predicate.[26]

(76) a. $\lambda x.g(x) \succeq \mathbf{s}(g)$ $\qquad g \in \{[\![\text{relative A}]\!]\}$

[26] I say 'at least' because there are classes of gradable predicates that I have not yet considered, such as color terms, 'extreme' adjectives like *horrified* and *gigantic* (see Paradis 2001), and 'evaluative' gradable adjectives like *beautiful* and *fun* (see Bierwisch 1989). Whether the truth conditions of the positive forms of these adjectives involve meanings like those in (76) or whether they require additional senses is an important question, but one that goes beyond the scope of this paper.

b. $\lambda x.g(x) \succ min(\text{SCALE}(g))$ $\quad g \in \{[\![\text{absolute min A}]\!]\}$
c. $\lambda x.g(x) = max(\text{SCALE}(g))$ $\quad g \in \{[\![\text{absolute max A}]\!]\}$

The theoretical questions we are faced with are the following. First, what (if anything) is the 'core meaning' of the positive form? Second, how can we explain (in a principled way) why particular adjectives are assigned particular interpretations in the positive form? In other words, what is the underlying explanation of the distinction between relative and absolute adjectives on the one hand, and between maximum absolutes and minimum absolutes on the other hand?

The answer to the first question emerges when we push the descriptive observations a bit further. Since relative and absolute gradable adjectives are of the same semantic type — both denote functions of type $\langle e, d\rangle$ — the variability illustrated in (76) must reflect variability in the interpretation of the positive morpheme (or positive form type-shifting rule; see the discussion in section 2.2), not in the interpretations of the adjectives.[27] In other words, *pos* is polysemous, having (at least) the three truth-conditionally distinct interpretations spelled out in (77).

(77) $[\![[_{\text{Deg}}\ pos]]\!] =$
a. $\lambda g \lambda x.g(x) \succeq \mathbf{s}(g)$
b. $\lambda g \lambda x.g(x) \succ min(\text{SCALE}(g))$
c. $\lambda g \lambda x.g(x) = max(\text{SCALE}(g))$

As these representations make clear, the various interpretations of *pos* share a fundamental feature: they turn a gradable adjective (a measure function) into a property of individuals by relating the degree to which an individual possesses the property measured by the adjective to a reference point that is computed as a function of the meaning of the adjective (the standard of comparison). The polysemy of the positive form arises from the fact that there are several natural options for establishing such a reference point. The maximal and minimal degrees on a scale are clearly two such options, since they are fixed values that

[27] That both relative and absolute adjectives denote functions of type $\langle e, d\rangle$ (or whatever type we assign to gradable adjectives; see note 5) is shown by the fact that both combine with comparative and other degree morphology; see the discussion of this point in section 3.1. We could avoid the assumption that *pos* is polysemous by assuming that absolute adjectives have both type $\langle e, d\rangle$ and type $\langle e, t\rangle$ interpretations, with the latter being whichever of (76b-c) is appropriate for the adjective. However, this approach would need to provide some mechanism to ensure that combination of *pos* with the type $\langle e, d\rangle$ form of an absolute adjective is blocked, since this would freely give rise to relative interpretations. If the proposals made below to regulate the interpretation of *pos* can apply equally to regulate the choice between $\langle e, d\rangle$ and $\langle e, t\rangle$ interpretations of absolute adjectives, then this hypothesis would be an option to consider.

can be computed strictly on the basis of the function expressed by the adjective (its range); a contextually significant degree of the property expressed by the adjective is a third, since it supports a meaningful partition of the domain of the adjective. Further investigation may reveal even more subtle shades of meaning for the positive form, but it is clear that at least the three options that have been the focus of this paper are attested.

This explanation of why *pos* has the range of interpretations it does still leaves open the question of why the positive form of a particular adjective is not itself polysemous — why the value of $[\![\text{pos}]\!]([\![\text{A}]\!])$ is fixed for particular choices of A. Why does *long* combine only with (77a), deriving a relative interpretation, while *bent* combines only with (77b) and *straight* only with (77c), deriving minimum and maximum absolute interpretations, respectively? In particular, why don't absolute adjectives like *straight* and *bent* also combine with (77a), deriving relative interpretations? Put another way, what factors determine that relative adjectives give rise to vague interpretations in the positive form, but absolute adjectives do not?

In the remainder of this paper, I will develop an answer to this question by arguing that the interpretation of the positive form — the choice of a particular value of *pos* — is determined by two factors: the structure of the scale used by an adjective and a general principle of interpretive economy, which requires the denotation of a constituent to be computed on the basis of the conventional meanings of its subconstituents to the extent possible, allowing for context sensitive truth conditions only as a last resort.

4.2. Scale structure and standard of comparison

The hypothesis that gradable adjectives denote functions that map objects onto representations of the degree to which they posses some gradable property (or that they incorporate such functions as part of their meanings; see note 5) leads to the expectation that gradable adjectives may differ with respect to features of those representations, i.e., their scales. There are a number of different ways that adjectival scales can be formalized, but minimally they must be triples $\langle P, \prec, \delta \rangle$ where P is a dense set of points, $\prec$ is a total ordering on P, and δ is a dimension (see Bartsch and Vennemann 1972, 1973; Bierwisch 1989; Kennedy 1999.

The dimension indicates the kind of property measured by the scale, and is the most obvious parameter of scalar variation, since it both distinguishes different adjectives from each other (e.g. *expensive* measures an object according to a dimension of COST while *fast* measures

an object according to a dimension of SPEED), and different senses of individual adjectives from each other (e.g., *long* can measure an object either with respect to LINEAR EXTENT or TEMPORAL EXTENT; see the discussion of indeterminacy in note 4).

The scalar feature that is relevant to the current discussion is the structure of the set P, which provides the range of the function expressed by the adjective: the set of degrees. Kennedy and McNally (2005) argue that gradable adjectives may differ with respect to the structure of this set, in particular, whether it does or does not have minimal or maximal elements, i.e., whether the scale is open or closed (see also Kennedy and McNally 1999; Paradis 2001; Rotstein and Winter 2004). There are four obvious ways in which scales could vary according to this feature, which are schematically represented in (78): a scale could lack either minimal or maximal elements, it could include a minimum but no maximum, it could include a maximum but no minimum, or it could include both a minimum and a maximum.[28]

(78) *A typology of scale strucures*

a. (TOTALLY) OPEN: ○——————————○

b. LOWER CLOSED: ●——————————○

c. UPPER CLOSED: ○——————————●

d. (TOTALLY) CLOSED: ●——————————●

As shown independently by Rotstein and Winter (2004) and Kennedy and McNally (2005), evidence for this typology comes from the distribution of adjectival modifiers that pick out maximal and minimal degrees on a scale. The reasoning relies on the following observation (which goes back at least to Sapir 1944): antonymous pairs of gradable adjectives map their arguments onto the same scale (*tall* and *short* both measure degree of height, *wet* and *dry* both measure amount of moisture, *full* and *empty* both measure level of contents, and so forth), but impose inverse orderings on their shared domains. This is reflected by equivalences like (79).

(79) a is {taller, fuller, wetter} than $b \Leftrightarrow$
b is {shorter, emptier, dryer} than a.

[28] A related but distinct feature that could be linguistically significant is whether a scale is bounded or unbounded. All closed scales are bounded on the relevant endpoint(s), but open scales may be further distinguished by whether they approach a value (e.g. 0) but do not include it, or whether they are completely unbounded. The representations in (i) are meant to abstract away from this distinction, so that ○ could in principle mean either 'open and bounded' or 'open and unbounded'.

Furthermore, positive adjectives like *tall*, *full* and *wet* measure increasing amounts of a property (if a is taller than b, then a has more height than b), while negative adjectives measure decreasing amounts of a property (if a is shorter than b, then a has less height than b).

There are different ways to formally capture this relation between polar antonyms, which depend on particular assumptions about the representation of degrees.[29] What is important for the current discussion is that the relation between polar antonyms and their shared scales leads to the following predictions if the scalar typology in (78) is real. First, if a scale is closed on the lower end, then the range of the positive member of an antonym pair that uses that scale should include a minimum degree (one that is ordered below all others) and the range of the negative member should include a maximum degree (one that is ordered above all others). Second, if a scale is closed on the upper end, then the range of the postive member of an antonym pair should include a maximum degree and the range of the negative member should include a minimum degree. Conversely, if a scale is open on the lower end, then the positive antonym should have no minimum degree and the negative antonym should have no maximum; and if a scale is open on the upper end, the positive should have no maximum and the negative no minimum.

The empirical probe that Rotstein and Winter 2004 and Kennedy and McNally 2005 use to test these predictions is acceptability of degree modifiers that pick out maximal or minimal degrees on the scales of the adjectives they modify. *Absolutely, completely, totally* and *perfectly* are examples of the former kind of degree modifier; *slightly* and *partially* are examples of the latter type. The general set of predictions for combinations of these modifiers with positive and negative adjectives using the four scale types illustrated above are laid out in (80).

[29] In particular, the crucial issue is whether they are characterized as points or intervals. On the former view, degrees can be identified directly with the points in P, and polarity is captured by assuming that antonyms differ with respect to the basic ordering on the scale: if e.g. *tall* denotes a function into $\langle P, \prec, \text{HEIGHT}\rangle$, then *short* denotes a function into $\langle P, \succ, \text{HEIGHT}\rangle$. On the latter view, degrees correspond to intervals based on P, and polarity is represented as a distinction in the ranges of antonyms: *tall* maps its arguments onto intervals that originate at the lower end of the scale; *short* maps its arguments onto intervals that originate at the upper end of the scale. For the purposes of this paper, it does not matter which model of degrees and polarity we adopt since the relation between maximal/minimal degrees and the structure of the sdcale that is exploited here is the same on either approach. See Kennedy 2001 for discussion of empirical issues that bear on this choice.

(80)

	OPEN	L-CLOSED	U-CLOSED	CLOSED
	$\text{Deg}_{max/min}$	$\text{Deg}_{max/min}$	$\text{Deg}_{max/min}$	$\text{Deg}_{max/min}$
A_{pos}	??/??	??/√	√/??	√/√
A_{neg}	??/??	√/??	??/√	√/√

The empirical picture is complicated somewhat by the fact that not all modifiers cooccur with all adjectives for apparently idiosyncratic reasons, and some of the maximizers can also be assigned 'high degree' rather than strictly maximal interpretations (see Kennedy and McNally 2005 for discussion of how to control for this). However, the modifiers *perfectly* and *slightly* provide clear judgments across a broad number of cases, and as shown by the examples in (81), demonstrate that the expected pattern does in fact emerge.

(81) *Open scales*

a. ??perfectly/??slightly {tall, deep, expensive, likely}
b. ??perfectly/??slightly {short, shallow, inexpensive, unlikely}

(82) *Lower closed scales*

a. ??perfectly/slightly {bent, bumpy, dirty, worried}
b. perfectly/??slightly {straight, flat, clean, unworried}

(83) *Upper closed scales*

a. perfectly/??slightly {certain, safe, pure, accurate}
b. ??perfectly/slightly {uncertain, dangerous, impure, inaccurate}

(84) *Closed scales*

a. perfectly/slightly {full, open, opaque}
b. perfectly/slightly {empty, closed, transparent}

There are a number of questions that should be asked about these patterns, most important of which is: why do the scales used by particular adjectives have the structure they do? For example, naive intuition suggests that the COST scale should have a minimal value representing complete lack of cost, just as the DIRT scale has a minimal value representing complete lack of dirt. However, the unacceptability of *??slightly/??partially expensinve* and *??perfectly/??completely/??absolutely inexpensive* (cf. *slightly/partially dirty* and *perfectly/completely/absolutely clean*) indicates that as far as the gradable adjective pair *expensive/inexpensive* is concerned, this is not the case: the scale used by these adjectives to represent measures of cost does not have a minimal element. Instead, the property of zero cost is named by the non-gradable adjective *free*. It is likely that the structure of a scale is largely determined by the nature

of the property that it is used to measure, but the different behavior of e.g. *expensive/inexpensive* vs. *dirty/clean* suggests that language may diverge from naive intuition.

Whatever principles explain why particular adjectival scales have the structures they do (see Kennedy and McNally 2005 for specific proposals regarding the scale structure of adjectives derived from verbs), what is important for the current discussion is that the scalar typology illustrated in (78) has consequences for the truth conditions of the positive form. As observed by Kennedy and McNally (2005) and as illustrated by the examples in (81)-(84), gradable adjectives that use totally open scales have relative interpretations in the postive form; gradable adjectives that use totally or partially closed scales have absolute interpretations. (Rotstein and Winter (2004) also observe the correlation between closed scales and absolute interpretations, but do not discuss the open scale/relative correlation.)

The first generalization is exceptionless, and follows directly from the interaction of scale structure and the semantics of *pos* proposed in (77). Open scales by definition lack maximal and minimal values, therefore combination of an open scale adjective such as *long* with either the minimum or maximum standard meanings of *pos* in (77b) and (77c) results in undefined truth conditions. The only remaining option is the 'significant degree' meaning in (77a), deriving a context dependent, relative interpretation of the positive form.

The second generalization is not exceptionless: as we have seen, relative interpretations of closed-scale adjectives can be forced in some contexts (e.g., by adding a *for*-PP; see the discussion of this point in section 3.2.3). But the full range of data discussed in section 3 strongly supports the conclusion that the standards of comparison for adjectives whose scales are closed on one end or the other default to an endopint of the scale: the minimum in some cases (*dirty*, *bent*, etc.); the maximum in others (*clean*, *straight*, etc.).

However, this generalization does not follow from the interaction of scale structure and the proposed semantics of *pos* alone. A closed scale structure may restrict the range of interpretations that *pos* can have (for example, the absolute maximum interpretation is incompatible with a lower closed adjective like *bent* for the same reason that it is blocked with an open scale adjective like *long*), but it does not uniquely determine a particular interpretation. Specifically, there is no semantic incompability between (totally or partially) closed scales and a relative interpretation of the positive form. The possibility of forcing such interpretations in specific contexts illustrates this, but at the same time, the robust and systematic evidence that adjectives that express functions to closed scales default to absolute interpretations

indicates that some other constraint, sensitive to this particular feature of gradable adjective meaning, is constraining the interpretation of the positive form.

4.3. Interpretive economy

In fact, we have already seen evidence that features of the conventional meaning of a gradable adjective can have systematic consequences for the truth conditions associated with the positive form. Recall from section 2.4 that when the domain of a relative adjective is explicitly restricted (by a *for*-PP), the restricted domain must be used as the basis for computing the standard of comparison; other properties in the context, no matter how salient, are ignored. This restriction was expressed in the descriptive constraint repeated in (85).

(85) *Domain Dependence*
If a gradable adjective has a restricted domain, then the standard of comparison must correspond to a significant degree of the relevant property relative to the domain.

The generalization about the relation between closed scales and absolute standards is similar in form (though it would more accruately labeled *Range Dependence*, given the analysis of gradable adjectives as functions from objects to degrees): if a gradable adjective expresses a function to a closed scale, then the standard of comparison must correspond to a maximal or minimal element of the scale.

I would like to suggest that the similarity between these two descriptive generalizations is not accidental, but rather that both follow from a more general principle of Interpretive Economy, which requires truth conditions to be computed on the basis of the conventional meanings of the expressions of a sentence (or logical form) to the extent possible, allowing for use of contextual information in the calculation of truth conditions only when conventional meaning is insufficient. This constraint is stated in the form of the directive in (86).

(86) *Interpretive Economy*
Maximize the contribution of the conventional meanings of the elements of a constituent to the comptuation of its meaning.

An important question is what kind of constraint (86) is. The fact that it can be overridden in certain contexts indicates that it is not a strict rule of grammar; at the same time, the fact that it affects entailments indicates that it is active at the level of semantic composition. The functional basis for *Interpretive Economy* is clear: although participants in a discourse may not be in full agreement about those properties of the

context that play a role in the computation of context-dependent features of meaning, they are in agreement about the conventional meanings of the words and complex expressions in the sentences they use to communicate (assuming they share the same lexicons and grammars). (86) maximizes the role of 'agreed upon meanings' in the computation of truth conditions, simplifying composition and constraining the range of interpretations that a particular structure may be assigned.

Given these considerations, I will assume for the purpose of this paper that *Interpretive Economy* is a constraint on semantic processing that applies locally at each instance of composition. This is clearly a hypothesis that should be further tested and developed in an experimental context; here I will focus on providing empirical support for it by showing the role it plays in explaining the facts discussed in this paper.

Consider first the case of open scale adjectives with unrestricted vs. restricted domains, as in (87a-b).

(87) a. That animal is a large mammal.
b. That animal is large for a mammal.

Since *large* is an open scale adjective, only the relative meaning of *pos* is an option for the interpretation of the positive form, which means that the adjectives (or rather the Degree Phrases they head) in (87a-b) express the properties in (88a-b). (I am assuming for the purpose of this example that the domain in (87a) is unrestricted, though it could be implicitly restricted.)

(88) a. $\lambda x.\mathbf{large}(x) \succeq \mathbf{s}(\lambda y.\mathbf{large}(y))$
b. $\lambda x : \mathbf{mammal}(x).\mathbf{large}(x) \succeq \mathbf{s}(\lambda y : \mathbf{mammal}(y).\mathbf{large}(y))$

The λ-expressions in (88a-b) are metalanguage representations of the truth conditions associated with these properties; in particular, $\mathbf{s}(g)$ is an abbreviation for 'the minimal degree that represents a significant amount of the property measured by g'. Let us assume that significance is an inherently relative notion: in order to determine the minimal degree that represents a significant amount of size, for example, you need a property on which to base this judgment. In the case of (88a), the conventional meaning of the adjective is a function from objects that have sizes to sizes.[30] However, the mere property of 'having size'

[30] That is, even an 'unrestricted' gradable adjective comes with an inherent domain restriction, as shown by their selectional restrictions. For example, the domain of *wide* is restricted to objects that can have a horizontal spatial dimension, explaining the anomaly of (non-metaphorical uses of) *??wide idea, ??wide claim, ??wide joke* and so forth.

is insufficient to determine what counts as a significant degree of size, so some other property must be recovered from the context. The property expressed by the modified nominal is highly salient and so is the most likely candidate, though as we have seen this is only a preference, not a requirement.

In contrast, the conventional meaning of the (modified) adjective in (88b) is a function from mammals (that have size) to sizes. The property of being a mammal provides a basis on which to calculate a significant degree of size; since this property is part of the conventional meaning of the adjective, *Interpretive Economy* dictates that it must be used in computing the truth conditions of the positive form. Use of a contextual property instead (or in addition) to calculate significance, no matter how salient, is prohibited. Note that context sensitivity is not entirely eliminated even in this case, since what counts as significant is also a function of the interests and expectations of the participants in the discourse (Graff 2000). However, use of the restricted domain maximizes the contribution of the conventional meaning of the constituents of the predicate to the calculation of its truth conditions, in accord with (86).

Turning now to closed scale adjectives, the effect of (86) is to choose between the competing senses of *pos*, favoring an absolute interpretation over the relative interpretation. Consider for example the case of *bent/straight*, which use a lower-closed scale according to the diagnostics discussed in the previous section. For the positive adjective *bent*, both of the interpretations in (89) are in principle possible.

(89) a. $\lambda x.\mathbf{bent}(x) \succeq \mathbf{s}(\lambda y.\mathbf{bent}(y))$
b. $\lambda x.\mathbf{bent}(x) \succ min(\textsc{scale}(\lambda y.\mathbf{bent}(y)))$

However, the truth conditions of (89b) are computed strictly on the basis of the conventional meaning of *bent* (its scale), while (89a) introduces context dependence in the significant degree computation. *Interpretive Economy* therefore dictates that (89b) should be preferred to (89a), deriving a minimum standard interpretation.

The analysis of the positive form of *straight* is essentially the same, except that the choice is between the relative interpretation in (90a) and the maximum standard intepretation in (90b).

(90) a. $\lambda x.\mathbf{straight}(x) \succeq \mathbf{s}(\lambda y.\mathbf{straight}(y))$
b. $\lambda x.\mathbf{straight}(x) = max(\textsc{scale}(\lambda y.\mathbf{straight}(y)))$

Again, the absolute interpretation is favored because it allows the meaning of the predicate to be computed strictly on the basis of the conventional meanings of its constituents.

Interpretive Economy makes clear predictions about gradable adjectives with lower and upper closed scales, since the only competing interpretations are the relative one and one of the absolute ones (whichever is appropriate given the scale structure of the adjective). The case of adjectives with totally closed scales is more complicated, however. Such adjectives are correctly predicted to be incompatible with relative interpetations, since *Interpretive Economy* always favors an absolute interpretation over a relative one if the former is an option. However, the scale structure of such adjectives is compatible with either of the two absolute interpretations (minimum standard or maximum standard). This suggests that such adjectives should display interpretive variability in the positive form, taking on maximum standard interpretations in some contexts and minimum standard interpretations in others.

The antonyms *opaque* and *transparent* verify this prediction. According to the diagnostics discused above, these use a totally closed scale (*completely/slightly opaque/transparent*), and so are in principle compatible with either minimum or maximum standard intepretations in the positive form. The following examples show that both interpretations are in principle possible. Consider a context in which I am manipulating a device that changes the degree of tint of a car window from 0% (completely transparent) to 100% (completely opaque). (91a) can be felicitously utterered at the point at which I have almost reached 100% of tint, demonstrating both that *opaque* can have a maximum standard (I am denying that the glass is completely opaque) and that *transparent* can have a minimum standard (partial transparency).

(91) a. The glass is almost opaque, but not quite. It's still transparent.
b. The glass is almost transparent, but not quite. It's still opaque.

Likewise, (91b) can be used to describe the reverse situation: one in which I have dialed down almost to 0% of tint. Here *transparent* has a maximum standard (comoplete transparency) and *opaque* has a minimum standard (partial opacity).

Out of context, there is a preference for maximum standard interpretations of *opaque* and *transparent*, but this can be explained in pragmatic terms: for any closed scale adjective, a maximum standard interpretation entails a minimum one, but not vice-versa. Assuming that stronger meanings are in general favored (see e.g. PETERS ET AL's analysis of the interpretive variability of reciprocal constructions), this preference follows. If miminum standard interpretations

were impossible, however, then the second sentences in (91a-b) would be contradictory.

Other total scale adjectives that allow both minumum and maximum standard intepretations are *open, exposed* and *uncovered*. The examples in (92) illustrate this for *open*.[31]

(92) a. If the airlock is open, the cabin will depressurize.
b. The ship can't be taken out of the station until the space door is open.

If I am a member of the crew of the starship *Enterprise* and I do not understand (92a) to be a warning that any amount of opening of the airlock will result in depressurization, then I am a danger to the ship and crew. Likewise, if I am the helmsman and fail to understand (92b) as an prohibition against trying to leave the station before the space door is completely open (here *the space door* refers to the door of the space station, which the ship needs to pass through in order to get into space), I am again a danger to the ship and crew.

Despite these examples, however, many (if not most) totally closed scale adjectives have fixed standards of comparison. For example, *closed, hidden* and *covered*, the antonyms of *open, exposed* and *uncovered*, all have maximum standards; so do the pair *full/empty*. Others, such as the deverbal adjectives *acquainted (with y)* and *documented* have minimum standards. (These claims can be verified by examining the adjectives' entailment patterns; see section 3.2.2.) It is possible that these adjectives are somehow lexically marked to combine only with the appropriate version of *pos*, but this would raise the question of why these particular total scale adjectives are restricted in this way, while the ones discussed above are not.

In fact, closer inspection of totally closed scale adjectives reveals a systematic relation between features of their meaning and choice of absolute standard. The picture is clearest if we focus on the case of deverbal adjectives. As pointed out by Kennedy and McNally (2005), there is a correlation between the scale structure of a deverbal adjective and the event structure of its source verb: if the verb is telic, its adjectival form has a totally closed scale; if it is atelic, it has a partially closed scale. Kennedy and McNally further show that in the specific case of total scale deverbal adjectives derived from accomplishment verbs,

[31] Similar examples can be constructed for *exposed* and *uncovered*. For example, *an exposed line of troops* is most naturally understood to describe a line of troops in which some soldiers are exposed to enemy fire (as pointed out to me by Mark Richard). However, the uproar over *Janet Jackson's exposed breast* during the halftime show of Superbowl XXXVIII resulted not because there was some visible skin, but because her breast was completely exposed, i.e., exposed to the nipple.

the orientation of the absolute standard (minimum vs. maximum) is a function of the role played by the argument of the adjective in the event described by the verb: if the argument corresponds to an incremental theme, then the adjectival form has a maximum standard; if it does not, then the adjectival form has a minimum standard.

This difference is illustrated quite clearly by the adjectival form(s) of *loaded*. When the verb is turned into an adjective that measures the degree to which its goal argument is loaded with some contents (*loaded with*), it has a minimum standard, but when it is turned into an adjective that measures the degree to which its incremental theme argument is loaded into a container (*loaded on*), it has a maximum standard. (See Dowty 1991 for arguments that the 'contents' argument of *load* is the incremental theme.) This is illustrated by the contrast in (93): (93a) is consistent with a non-maximal degree of 'loaded-with-ness'; (93b) is inconsistent with a non-maximal degree of 'loaded-on-ness'.

(93) a. The truck is loaded with the boxes, but half of it remains empty.
b. ??The boxes are loaded on the truck, but half of them are still on the dock.

Kennedy and McNally's explanation for this correlation is as follows. Assume that the positive form of a deverbal adjective entails the completion of an eventuality corresponding to the one described by the source verb. In the case of an adjective whose argument is an incremental theme, the relevant event is not completed unless the argument has been totally affected by the verb; this will only be the case if it has a maximum degree of the property expressed by the adjective. In the case of a non-incremental theme argument, this relation does not hold; instead, the completion of the event is consistent with a situation in which the argument merely has a non-zero degree of the property expressed by the adjective. The prediction is that deverbal adjectives with incremental theme arguments will combine only with the maximum standard interpretation of *pos*, since the minimum standard interpretation is incompatible with the 'completed event entailment'. In contrast, adjectives with non-incremental theme arguments should in principle be able to combine with either version, and so should show variability in their truth conditions.

These predictions are verified by further examination of *loaded*. (93b) showed that it has only a maximum standard interpretation with an incremental theme argument. (If it also allowed a minimum standard interpretation, then (93b) would not be a contradiction.) (93a) showed

that *loaded* can have a minimum standard when predicated of a non-incremental theme arguement; (94) shows that it can also be assigned a maximum standard.

(94) All of the boxes are on the truck, but it's not loaded yet. We still need to get the furniture in there.

Here what is being denied is that the truck is fully loaded, not that it has some degree of loadedness; this is a maxumum standard interpretation.

Most adjectives with totally closed scales appear to be deverbal (perhaps because such scales are so closely tied to event structures, as argued by Kennedy and McNally (2005)), though not all are related to accomplishment verbs. The pairs *open/closed, covered/uncovered, hidden/exposed* are related to achievements, but the core of the explanation outlined above carries over directly. The difference between the members of such adjectival pairs is that the verbal source of one member (*open, uncover, expose*) names the initiation point of an event, while the verbal source of the other (*close, cover, hide*) names the culmination point. As a result, an argument of the 'initiation predicate' will count as having participated in the event named by the verb just in case it has a minimal degree of the measured property, while an argument of the 'culimation predicate' will count as having participated in the event just in case it has a maximal degree of the relevant property. The latter are therefore predicted to have maximal standards, and the former to have variable standards, as we have seen.

Finally, the difference between the lexical adjectives *opaque/transparent* on the one hand and *full/empty* can be explained if we assume that the latter but not the former are related to events in the same way as deverbal adjectives. One obvious argument in favor of this position is the fact that the arguments of the latter pair are incremental themes of near-identical verbal counterparts (*full/empty*), while the former do not even have verbal counterparts (*??opacify/??transparentize*). Another is the fact that *full/empty* are stage level (in the positive form), while *opaque/transparent* are individual level, as shown by the difference in interpretation of bare plural subjects:

(95) a. Cisterns are full/empty. EXISTENTIAL
b. High-rise windows are opaque/transparent. GENERIC

Assuming the stage/individual contrast correlates with the presence of an event variable in the predicate, this contrast suggests that *full/empty* are eventive in a way that *opaque/transparent* are not. Their interpretations in the positive form then follow from the considerations out-

lined above. The positive form of *full/empty* introduces an entailment that their arguments have (maximally) participated in an event of filling/emptying; this rules out a minimum standard interpretation. The positive form of *opaque/transparent* has no such entailment, so both the maximum and minimum standard interpretations are in principle possible, though pragmatic considerations should favor the former.

4.4. Vagueness and gradability

The assumptions about gradable adjective meaning that I adopted at the beginning of this paper, in which gradable adjectives definie mappings between objects and abstract reprsentations of measurement (scales and degrees), which in turn have crucial representational and ontological status in the semantics, represent one of two main approaches to the semantics of gradable predicates. The other kind of approach, developed in most detail as a compositional theory of gradable predicate meaning by Klein (1980) (see also McConnell-Ginet 1973; Fine 1975; Kamp 1975; Larson 1988), does not introduce the semantic type 'degree' into the semantics, but instead analyzes gradable predicates in a way parallel to other predicative exressions, as functions of type $\langle e, t\rangle$. What is special about gradable predicates is that they denote context-sensitive, partial functions from individuals to truth values: in addition to their positive and negative extensions, they may have an 'extension gap', which corresponds to the set of objects that the predicate is neither true nor false of in a particular context of utterance.

Crucially, the positive and negative extensions and the extension gap of a gradable predicate may vary across contexts of use, becoming more or less precise.[32] This variability underlies the properties of vagueness that have been the focus of this paper, and provides the basis for a semantics of comparatives and other complex forms. Comparatives, for example, involve quantification over possible interpretations (or 'precisifications', to use Pinkal's (1995) terminology) of an adjective: *x is more A than y* is true just in case there is an interpretation of A such that that x is in its positive extension but y is not. In effect, this type of apporach derives gradability from a general semantics for vague predicates, while degree-based approaches build an account of vagueness on top of a more general semantics of gradability.

[32] However, the contextual variability of the function expressed by the base adjective must also be subject to general constraints that reproduce the effects of a scalar semantics. For example, if an object x is in the positive extension of A in context c, and $x \succ y$ with respect to the property described by A, then there is no context c' such that y is in its positive extension of A in c' but x is not; see e.g. Klein's (1980) *Consistency Postulate*.

A potential advantage of the 'vague predicate' analysis is that it treats the positive form as basic: there is no need for a special positive form morpheme or type shifting rule, as in degree-based analyses, and morphological (un)markedness (of e.g. the positive vs. comparative forms) is mirrored by compositonal complexity. However, the general question that I have attempted to answer here — what factors determine the truth conditions associated with the positive form of a particular gradable adjective in a particular context? — is just as relevant for this type of approach as it is for degree-based approaches. In particular, an analysis that derives gradability from a general, non-scalar semantics for vague predicates must explain the empirical phenomena that have been the focus of this paper: the semantic properties of relative and absolute gradable adjectives in the positive form. While it may be difficult but not impossible to explain some of these features (e.g., *Domain Dependence* and the impossibility of crisp judgments for relative adjectives in the positive form), I do not see how such an approach can account for the basic facts of the relative/absolute distinction. Since vagueness (i.e., allowing for variable interpretations/precisifications) is a necessary condition for comparison, the expectation is essentially the same as the 'traditional' degree-based semantics of the positive form: all gradable predicates should be vague. The challenge for a non-degree based analysis is to explain why only relative adjectives are vague in the positive form, while absolute adjectives have fixed positive and negative extensions, but remain fully gradable.

In this paper, I have argued that the explanation for this difference in interpretation is based on the structures of the scales onto which relative and absolute gradable adjectives map their arguments. In a degree-based semantics of gradable adjectives, scale structure is a basic component of the conventional meaning of a gradable adjective: it is a property of the range of the measure function expressed by the adjective. As such, it may be explicitly or implicitly mentioned in the semantics of other expressions (such as the absolute maximum/minimum senses of the positive form morpheme/type shifting rule or other types of degree morphology; see Kennedy and McNally 2005), and it plays a crucial role in the evaluation of *Interpretive Economy*, effectively restricting the range of interpretations that a closed-scale adjective can have. The vagueness of the positive form is thus a function both of compositional semantics (via the meaning of *pos*) and lexical semantics (the scale structure of a gradable adjective), mediated by *Interpretive Economy*. It is in this sense that vagueness is directly influenced by the grammatical properties of the expressions that (potentially) give rise to it.

One way that a non-degree approach could adopt this explanation of the relative/absolute distinction would be to define the truth conditions of a gradable predicate in terms of scales and degrees, without actually introducing degrees into the inventory of semantic types, as in Barker 2002. Instead of a positive form morpheme/rule that maps a measure function to an appropriate property of individuals, such an approach would need to define a function from contexts to predicate denotations that makes the right use of scalar information when fixing the positive and negative extensions of a predicate, so that e.g. an adjective whose truth conditions are defined in terms of a lower-closed scale ends up denoting a property of individuals that is true just of an object just in case it has a non-zero degree of the relevant gradable property, and so forth. Such an analysis would be functionally equivalent to the one proposed here, and would rely on exactly the same principles to capture the relative/absolute distinction. More generally, such an approach would require just as strong a commitment to the linguistic significance of scale structure, further indicating the importance of scalar representations in natural language semantics.

5. Conclusion

This paper has developed a semantic analysis of gradable adjectives in the positive form in which the construction itself is polysemous, allowing in principle for at least three distinct interpretations: a relative one, a minimum standard absolute one, and a maximum standard absolute one. The actual truth conditions of the predicate, and in particular whether it gives rise to vagueness or not, are a function of the interaction of lexical properties of gradable adjectives (their scale structure, their domains) and a general constraint on *Interpretive Economy* that requires maximization of the contribution of the conventional meanings of the subparts of a constituent to the computation of its truth conditions.

The specific analysis of the vagueness of relative gradable predicates that I have advocated here is one that is stems from the semantics of the (relative interpretation of) positive form: a relative predicate is true of an object if it possesses a significant amount of the property measured by the relative adjective, as argued by Graff (2000) and Bogusławski (1975). Whether this analysis extends to an analysis of vagueness more generally is an issue that must be addressed in future work. While it seems straightforward to apply the proposals here to the case of vague determiners like *many/few* and *much/little* (construed as measure functions on pluralities and substances, respectively; see Hackl 2000)

and possibly even scalar nouns like *heap, pile* and so forth, it is unlikely that the analysis will subsume all cases of vagueness. For example, it seems quite undesirable to explain the fact that even a relatively stable common noun like *table* can give rise to borderline cases (Is a rock with a flat top a table? It depends on how it is being used....) by assigning it a semantic anlaysis as an open-scale function from objects to degrees of 'tablehood', which combines with *pos* to derive a function from individuals to truth values that is true of an object if it possesses a significant degree of 'tablehood'.[33]

In the end, it may very well be the case that vagueness can arise in two ways: from the conventional meanings of certain expressions (such as the positive form of relative gradable adjectives) and from more general principles governing the mapping from linguistic expressions to truth conditions (such as those described in section 4.4). A conclusion of the work presented here, however, is that we cannot only assume the latter, even though that would be theoretically most perspicuous. If the relative/absolute distinction stems from the interaction scalar properties of gradable adjectives and *Interpretive Economy*, as I have argued, then we have at least one case in which conventional meaning alone is the basis for whether or not two expressions of the same semantic type and with essentially the same kind of meaning will give rise to vagueness.

References

Abney, S. P.: 1987, 'The English Noun Phrase in its Sentential Aspect'. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, Mass.

Barker, C.: 2002, 'The Dynamics of Vagueness'. *Linguistics and Philosophy* **25**(1), 1–36.

Bartsch, R. and T. Vennemann: 1972, 'The Grammar of Relative Adjectives and Comparison'. *Linguistische Berichte* **20**, 19–32.

[33] Clear evidence that this is the wrong semantic analysis comes from the fact that *table* is not gradable:

(i) *That rock is more table than this bush.

However, it is possible in English to turn a nominal like *table* into a measure function that has essentially the features described above: (ii) essentially means that the rock has more of the qualities appropriate for tablehood than the bush does.

(ii) That rock is more of a table than this bush.

The acceptability of (ii) shows that the grammar contains mechanisms to make non-gradable expressions gradable (in this example, the '*of a*-construction'); at the same time, the anomaly of (i) shows that not all expressions that give rise to vagueness should be treated as underlyingly gradable.

Bartsch, R. and T. Vennemann: 1973, *Semantic Structures: A Study in the Relation between Syntax and Semantics*. Frankfurt: Athäenum Verlag.

Bhatt, R. and R. Pancheva: 2004, 'Late Merger of Degree Clauses'. *Linguistic Inquiry* **35**, 1–46.

Bierwisch, M.: 1989, 'The Semantics of Gradation'. In: M. Bierwisch and E. Lang (eds.): *Dimensional Adjectives*. Berlin: Springer-Verlag, pp. 71–261.

Bogusławski, A.: 1975, 'Measures are Measures: In Defence of the Diversity of Comparatives and Positives'. *Linguistiche Berichte* **36**, 1–9.

Corver, N.: 1990, 'The Syntax of Left Branch Extractions'. Ph.D. thesis, Tilburg University.

Cresswell, M. J.: 1977, 'The Semantics of Degree'. In: B. Partee (ed.): *Montague Grammar*. New York: Academic Press, pp. 261–292.

Cruse, D. A.: 1986, *Lexical Semantics*. Cambridge, UK: Cambridge University Press.

Dowty, D. R.: 1991, 'Thematic Proto-Roles and Argument Selection'. *Language* **67**, 547–619.

Fine, K.: 1975, 'Vagueness, Truth and Logic'. *Synthese* **30**, 265–300.

Giannakidou, A.: 2004, 'Domain Restriction and the Arguments of Quantificational Determiners'. In: Editor (ed.): *Proceedings of SALT 14*. Ithaca, NY, pp. xx–xx.

Graff, D.: 2000, 'Shifting Sands: An Interest-Relative Theory of Vagueness'. *Philosophical Topics* **20**, 45–81.

Grimshaw, J.: 1991, 'Extended Projection'. Unpublished ms., Brandeis University.

Hackl, M.: 2000, 'Comparative Quantifiers'. Ph.D. thesis, Massachusetts Institute of Technology.

Hankamer, J.: 1973, 'Why there are two 'than's in English'. In: C. Corum, T. C. Smith-Stark, and A. Weiser (eds.): *Proceedings of the 9th Annual Meeting of the Chicago Linguistics Society*. Chicago, IL.

Heim, I.: 1985, 'Notes on Comparatives and Related Matters'. Ms., University of Texas.

Heim, I.: 2000, 'Degree Operators and Scope'. In: B. Jackson and T. Matthews (eds.): *Semantics and Linguistic Theory 10*. Ithaca, NY, pp. 40–64.

Heim, I. and A. Kratzer: 1998, *Semantics in Generative Grammar*. Oxford: Blackwell.

Hellan, L.: 1981, *Towards an Integrated Analysis of Comparatives*. Tübingen: Narr.

Hoeksema, J.: 1984, 'Negative Polarity and the Comparative'. *Natural Language & Linguistic Theory* **1**, 403–434.

Kamp, H.: 1975, 'Two Theories of Adjectives'. In: E. Keenan (ed.): *Formal Semantics of Natural Language*. Cambridge: Cambridge University Press, pp. 123–155.

Kennedy, C.: 1999, *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. New York: Garland. (1997 UCSC Ph.D thesis).

Kennedy, C.: 2001, 'Polar Opposition and the Ontology of 'Degrees''. *Linguistics and Philosophy* **24**, 33–70.

Kennedy, C.: 2005, 'The Grammar of Comparison'. Paper presented at the Università di Venezia, June 6-7, 2005.

Kennedy, C. and L. McNally: 1999, 'From Event Structure to Scale Structure: Degree Modification in Deverbal Adjectives'. In: T. Matthews and D. Strolovitch (eds.): *Semantics and Linguistic Theory 9*. Ithaca, NY, pp. 163–180.

Kennedy, C. and L. McNally: 2005, 'Scale Structure and the Semantic Typology of Gradable Predicates'. *Language* **81**(2), 1–37.

Klein, E.: 1980, 'A Semantics for Positive and Comparative Adjectives'. *Linguistics and Philosophy* **4**, 1–45.

Klein, E.: 1991, 'Comparatives'. In: A. von Stechow and D. Wunderlich (eds.): *Semantik: Ein Internationales Handbuch der Zeitgenössischen Forschung*. Berlin: Walter de Gruyter, pp. 673–691.

Kyburg, A. and M. Morreau: 2000, 'Fitting Words: Vague Language in Context'. *Linguistics and Philosophy* **23**(6), 577–597.

Larson, R. K.: 1988, 'Scope and comparatives'. *Linguistics and Philosophy* **11**, 1–26.

Lasersohn, P.: 1999, 'Pragmatic Halos'. *Language* **75**(3), 522–551.

Lewis, D. K.: 1970, 'General Semantics'. *Synthese* **22**, 18–67.

Ludlow, P.: 1989, 'Implicit Comparison Classes'. *Linguistics and Philosophy* **12**, 519–533.

Martí, M. L.: 2002, 'Contextual Variables'. Ph.D. thesis, University of Connecticut.

McConnell-Ginet, S.: 1973, 'Comparative Constructions in English: A Syntactic and Semantic Analysis'. Ph.D. thesis, University of Rochester.

Meier, C.: 2003, 'The Meaning of *Too*, *Enough*, and *So ... That*'. *Natural Language Semantics* **11**(1), 69–107.

Montague, R.: 1974, 'English as a Formal Language'. In: R. Thomason (ed.): *Formal Philosophy: Selected Papers*. New Haven: Yale University Press, pp. 188–221.

Neeleman, A., H. Van de Koot, and J. Doetjes: 2004, 'Degree Expressions'. *The Linguistic Review* **21**(21), 1–66.

Paradis, C.: 2001, 'Adjectives and Boundedness'. *Cognitive Linguistics* **12**(1), 47–65.

Pinkal, M.: 1979, 'Semantics from Different Points of View'. In: R. Bäurle, U. Egli, and A. von Stechow (eds.): *How to Refer with Vague Descriptions*. Berlin: Springer-Verlag, pp. 32–50.

Pinkal, M.: 1995, *Logic and Lexicon*. Dordrecht: Kluwer.

Rotstein, C. and Y. Winter: 2004, 'Total Adjectives vs. Partial Adjectives: Scale Structure and Higher-Order Modifiers'. *Natural Language Semantics* **12**, 259–288.

Rusiecki, J.: 1985, *On Adjectives and Comparison in English*. New York: Longman Linguistics Library.

Sapir, E.: 1944, 'Grading: A Study in Semantics'. *Philosophy of Science* **11**, 93–116.

Schwarzschild, R.: 2004, 'Measure Phrases as Modifiers of Adjectives'. To appear in *Recherches Linguistiques de Vincennes*.

Sedivy, J., M. Tanenhaus, C. Chambers, and G. Carlson: 1999, 'Achieving Incremental Semantic Interpretation through Contextual Representations'. *Cognition* **71**, 109–147.

Seuren, P. A.: 1973, 'The Comparative'. In: F. Kiefer and N. Ruwet (eds.): *Generative Grammar in Europe*. Dordrecht: Riedel, pp. 528–564.

Stanley, J.: 2000, 'Context and Logical Form'. *Linguistics and Philosophy* **23.4**, 391–434.

Stanley, J.: 2002, 'Nominal Restriction'. In: G. Peter and G. Preyer (eds.): *Logical Form and Language*. Oxford: Oxford University Press, pp. 365–368.

Stanley, J. and Z. Szabó: 2000, 'On Quantifier Domain Restriction'. *Mind and Language* **15**, 219–261.

Svenonious, P. and C. Kennedy: 2005, 'Northern Norwegian Degree Questions and the Grammar of Measurement'. Paper presented at the 31st Incontro di Grammatica Generativa, Universitá Roma 3, 25 February, 2005.

Sybesma, R.: 1999, *The Mandarin VP*. Boston: Kluwer.

Syrett, K., E. Bradley, C. Kennedy, and J. Lidz: 2005, 'Shifting Standards: Childrens Understanding of Gradable Adjectives'. Paper presented at the 2005 Annual Meeting of the Linguistics Socieity of America, San Francisco, CA.

Unger, P.: 1975, *Ignorance*. Oxford: Clarendon Press.

von Fintel, K.: 1994, 'Restrictions on Quantifier Domains'. Ph.D. thesis, University of Massachusetts, Amherst.
von Stechow, A.: 1984, 'Comparing Semantic Theories of Comparison'. *Journal of Semantics* **3**, 1–77.
Wheeler, S.: 1972, 'Attributives and Their Modifiers'. *Noûs* **6**(4), 310–334.
Williamson, T.: 1994, *Vagueness*. London: Routledge.
Yoon, Y.: 1996, 'Total and Partial Predicates and the Weak and Strong Interpretations'. *Natural Language Semantics* **4**, 217–236.

Contact address :

Department of Linguistics
University of Chicago
1010 E. 59th St.
Chicago, IL 60637 USA
ck@uchicago.edu