Towards a Grammar of Vagueness*

Christopher Kennedy
Northwestern University

DRAFT: April 5, 2004

This paper investigates the grammatical principles governing the interpretation of a class of vague predicates — gradable adjectives — focusing on the context dependence of the 'standard of comparison' with respect to which these predicates are judged to be true. I show that the range of variability in interpretation of the standard of comparison is broader than has generally been assumed, specifically that there are distinct types of standards with distinct effects on truth conditions, and I argue that an empirically adequate semantics for such predicates must be able to account for this variability. I then argue that the observed range of interpretations can be explained in terms of the interaction of the scalar properties of gradable predicates and general constraints on semantic economy and strength.

1 Introduction

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

(1) The Mars Pathfinder mission was expensive

The problem presented by sentences of this sort is that they are vague: (1) might be judged true in a conversation about things with the name 'Pathfinder' (such as compasses, mountain bikes, and sport utility vehicles, as well as missions to missions to Mars), but false in a discussion of the cost of various missions to outer space (since a notable fact about the Mars Pathfinder mission was its relatively low cost).

The standard analysis of these sorts of sentences is that the locus of vagueness is the gradable adjective *expensive*: gradable adjectives establish relations between objects and measures of the degree to which they possess some property, and predicates like *is expensive* denote the property of having a degree of cost that exceeds some standard of comparison of cost. Crucially, the value of the implicit standard is not specified in the lexical entry of *expen-*

*Acknowledgements.

sive, but is rather set contextually, and so it may vary in different utterance contexts. For example, in the first context mentioned above (things with the name 'Pathfinder'), the standard of comparison for *expensive* should be fairly low on the scale of cost, making (1) true. In the second context (missions to outer space), the standard of comparison should be considerably higher, making it possible for (1) to be false.

There is general agreement among researchers on this topic that something like this is what is going on in the interpretation of gradable adjectives. What is less clear, however, is how exactly the standard of comparison is actually determined in a context of use: to what extent is it determined compositionally by pieces of the linguistic representation (and how do these pieces interact), and to what extent do purely contextual (possibly non-linguistic) factors influence its value? In other words, what is the grammar of vagueness? This question has received a fair amount of discussion in the literature, though a completely satisfactory theory of the grammar of vagueness has not yet been achieved (see e.g., Sapir 1944; McConnell-Ginet 1973; Kamp 1975; Fine 1975; Klein 1980; Ludlow 1989; Kennedy 1999a; Graff 2000).

The goal of this paper is twofold. First, I will show that the empirical domain that such a theory needs to cover is actually quite a bit richer and more complicated than most of the research on this topic (which has focused almost exclusively on cases like (1)) assumes. Specifically, I will demonstrate that there are actually distinct types of standards of comparison, each of which have their own unique and clearly identifiable semantic properties but which also share enough features in common that we should expect a theory of the grammar of vagueness to provide a fully general account of them if it is to achieve explanatory adequacy.

Second, I will attempt to present such a theory, arguing that the facts can be explained in terms of the interaction of lexical properties of gradable predicates — in particular, the scalar structures they invoke as part of their semantics — and a general constraint on 'interpretive economy' that requires the values of free variable to be fixed based on the conventional meanings of expressions in a sentence whenever possible, allowing purely contextual determination of the values of such expressions only as a last resort.

2 The starting point

2.1 The semantics of gradable predicates

I begin with an overview of the semantic analysis of gradable adjectives. For concreteness, I will adopt the assumptions in (2), which are shared it some form

by many semantic analyses of gradable predicates (e.g., Seuren 1973; Bartsch and Vennemann 1973; Cresswell 1977; Hellan 1981; von Stechow 1984; Heim 1985, 2000; Bierwisch 1989; Klein 1991; Kennedy 1999a and others).[1]

(2) i. Gradable adjectives map their arguments onto abstract representations of measurement, or DEGREES.
ii. Degrees are formalized as points or intervals totally ordered along some DIMENSION (e.g., height, cost, etc.; the set of ordered degrees corresponds to a SCALE.
iii. Propositions constructed out of gradable adjectives express relations between degrees on a scale.

The most common implementation of this general view assumes that gradable adjectives denote relations between individuals and degrees. Specifically, GAs contain as part of their meanings a measure function and a partial ordering relation, as illustrated in (3) for the adjective *expensive*. (See Bartsch and Vennemann 1973; Kennedy 1999a for approaches in which the adjective just denotes a measure function.)

(3) a. $[\![[_{\text{A}}\ \text{expensive}]]\!] = \lambda d \lambda x.\mathbf{expensive}(x) \succeq d$
b. **expensive** = a function from objects to (positive) degrees of cost

This approach provides a straightforward analysis of GA+degree morpheme constructions like comparatives. Specifically, we can assume that degree morphemes have interpretations along the lines of (4): they impose restrictions on the semantic value of the degree argument of the adjective.

(4) $[\![\text{Deg(P)}]\!] = \lambda G \lambda x.\exists d[\mathbf{R}(d) \wedge G(d)(x)]$

Different degree morphemes assign different values to the restrictive clause **R**. For example, the comparative morphemes restrict the degree argument as shown in (5), where d_c is the semantic value of the comparative clause.

[1] These assumptions represent one of two main approaches to the semantics of gradable predicates. The other approach, developed in most detail by Klein (1980) (see also McConnell-Ginet 1973; Kamp 1975), does not introduce degrees into the ontology. Instead, this sort of approach analyzes gradable predicates as expressions of type $\langle e, t\rangle$ whose positive and negative extensions may vary across contexts of use (and must be determined in context, just like a standard of comparison). I believe that the general question I am addressing here — how do we settle on a particular set of truth conditions for a gradable predicate in a context of utterance? — is just as relevant for this sort of analysis as it is for degree-based approaches. I focus on the latter primarily for reasons of simplicity, though whether a Klein-style analysis can handle the facts I discuss below better than a degree-based approach is a question that should eventually be explicitly addressed.

(5) a. more: $\mathbf{R} = \lambda d.d \succ d_c$
b. less: $\mathbf{R} = \lambda d.d \prec d_c$
c. as: $\mathbf{R} = \lambda d.d \succeq d_c$

The details of compositional interpretation of comparatives vary depending on assumptions about syntax and the syntax-semantics interface (see e.g. Heim 1985, 2000; von Stechow 1984; Bierwisch 1989; Gawron 1995 for different implementations). One way to handle a sentence like (6a) is shown in (6b), where the comparative morpheme and comparative clause form a syntactic and semantic constituent, and the comparative clause contains an elided occurrence of the adjective *expensive* (indicated by struck-through text).

(6) a. The Viking Mission was more expensive than the Pathfinder Mission (was).
b. The Viking Mission was [$_{\text{AP}}$ [$_{\text{DegP}}$ -er [than *wh* the Pathfinder Mission ~~was *t* expensive~~]] expensive]

(7) a. $[\![[_{\text{A}}\ \text{expensive}]]\!] =$
$\lambda d\lambda x.\mathbf{expensive}(x) \succeq d$
b. $[\![\text{[than } wh \text{ the Pathfinder Mission ~~was } t \text{ expensive~~]}]\!] =$
$max\{d' \mid \mathbf{expensive}(p) \succeq d'\}$
c. $[\![[_{\text{DegP}}\ \text{-er [than } wh \text{ the Pathfinder Mission ~~was } t \text{ expensive~~]]}]\!] =$
$\lambda G\lambda x.\exists d[d \succ max\{d' \mid \mathbf{expensive}(p) \succeq d'\} \wedge G(d)(x)]$
d. $[\![[_{\text{AP}}\ [_{\text{DegP}}\ \text{-er than } wh \text{ the Pathfinder Mission ~~was } t \text{ expensive~~}]\ \text{expensive}]]\!] =$
$\lambda x.\exists d[d \succ max\{d' \mid \mathbf{expensive}(p) \succeq d'\} \wedge \mathbf{expensive}(x) \succeq d]$
e. $[\![(6a)]\!] = 1$ iff
$\exists d[d \succ max\{d' \mid \mathbf{expensive}(p) \succeq d'\} \wedge \mathbf{expensive}(v) \succeq d]$

2.2 The positive form

Paradoxically, it is the analysis of 'simple' predicates involving GAs — henceforth the POSITIVE FORM — that is the most problematic within this set of assumptions. The problem is that the analysis outlined above characterizes GAs as relational expressions, but in English at least, there is no overt degree morphology to saturate the degree argument of the adjective. The question, then, is how we get from an expression of type $\langle d, et\rangle$ to a property of individuals with the right truth conditions in the right contexts?

One possibility is to posit a null degree morpheme *pos*, as in Cresswell 1977; von Stechow 1984, which has the same semantic type as other degree morphology. Specifically, *pos* should have the denotation in (8), where $\mathbf{R}$ is

some appropriate property of degrees.

(8) $[\![[_{Deg}\ pos]]\!] = \lambda G_{\langle d,et\rangle}\lambda x.\exists d[\mathbf{R}(d) \wedge G(d)(x)]$

Potential evidence in favor of this approach comes from Mandarin Chinese. As shown by Sybesma (1999), the positive form of the adjective in Chinese must be morphologically marked by the morpheme *hen*, as shown in (9a).[2] The unmarked form in (9b) is actually ungrammatical if uttered in isolation, but in a context that supports a comparative interpretation, it is acceptable.

(9) 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 (10a-b), fully explicit comparative constructions are incompatible with *hen*.

(10) 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 the hypothesized null morpheme *pos*. While this does not provide conclusive evidence that English also has such a morpheme, it shows that such a morpheme is in fact an actual lexical item. Since positing a null version of *hen* for English allows us to have a fully compositional analysis of comparative and positive forms alike, I will make this assumption for the purposes of this paper.[3]

[2]The morpheme *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.

[3]An alternative to this approach would be to assume that when the degree argument of the adjective is not saturated by some linguistic expression, a default existential quantifier with a meaning along the lines of the one proposed below for *pos* is introduced, or that the degree argument is just saturated by a default free variable, whose value is fixed according to the semantic principles outlined below (see e.g. Barker 2002). Since the choice between *pos*, default quantification, or a free variable hinges on more general concerns about how implicit arguments are handled, rather than the specific proposals for the resolution of vagueness that I will make in this paper, I will assume the null morpheme analysis for convenience.

This approach to the positive form doesn't answer the question I started out with, however — how do we build determinate truth conditions for the positive form? — it just provides a formalism for restating it: what is the actual value of the degree restriction **R** in the positive form, and how does it allow for variation in truth conditions in different contexts of use? The answer to this question that is assumed by most analysis of gradable adjectives is that **R** is a context-depenedent property of degrees that is determined relative to some COMPARISON CLASS: a set of objects that are in some way similar to the subject of predication.

The basic idea can be illustrated by an example like (11), which is just like (1): it can be true or false in different contexts.

(11) Spartacus is old.

The intuition is that this sentence has different truth conditions depending on what we are comparing Spartacus to. In particular, if we know only that Spartacus is the name of a dog, then we might understand (11) to be equivalent to (12a). If, however, we know that Spartacus is the name of a pug dog, then we might understand (11) to be equivalent to (12b).

(12) a. Spartacus is old for a dog.
b. Spartacus is old for a pug dog.

In (12a-b), the *for*-PPs make explicit (two of) the implicit comparison classes that could be associated with (11). If we know that pugs tend to live longer than average dogs, then indeed (12a-b) have different truth conditions: (12b) requires Spartacus to be older than (12a) does.

Klein 1980 builds a detailed analysis of the positive form on top of these observations (albeit in a framework that makes different initial assumptions about the semantic type of GAs; see note 1). Roughly speaking, Klein argues that the domain of the predicate is restricted to just those objects that correspond to the comparison class under consideration, and that the 'cutoff point' for things that the predicate is definitely true of corresponds to a median value along an ordering of those objects.

In the context of the semantic analysis of gradable adjectives outlined in the previous section, we can implement Klein's idea by replacing the degree restriction variable **R** with a property of degrees that is defined in terms of an ordering with respect to some comparison class. In a context in which (11) is understood to mean (12a), this property is as in (13a); in a context in which (11) is understood as (12b), it has the value in (13b).[4]

[4] In her analysis of the Sorites Paradox, Graff (2000) argues that it is not enough to

(13) a. *λd.d is greater than a norm of age for the comparison class* DOGS
b. *λd.d is greater than a norm of age for the comparison class* PUGS

In other words, this analysis claims that contextual variability in gradable predicates does not involve full-scale variability of the restriction the degree quantifier (**R** in (8)), but rather variability in just a part of the restriction: the part that fixes the comparison class. The rest of the meaning — the requirement that the degree of the property possessed exceeds a norm — remains constant.

But what's a norm, and how do we incorporate the comparison classes in interpretations like (13a-b) compositionally? As an initial answer to the first question, let us assume that a norm is the mean of the degrees to which the things in the comparison class possess the relevant property. On this view, (13a-b) should be rewritten as (14a-b).[5]

(14) a. $\lambda d.d \succ mean\{d' \mid \exists x[x \in [\![\text{DOGS}]\!] \wedge \mathbf{old}(x) \succeq d']\}$
b. $\lambda d.d \succ mean\{d' \mid \exists x[x \in [\![\text{PUGS}]\!] \wedge \mathbf{old}(x) \succeq d']\}$

As for the second question, an important clue comes from the fact that sentences of the form *x is GrAdj for a NP* presuppose that *x* is an NP. All of

simply exceed the norm, as in (13), but rather a gradable predicate holds of an object only if that object exceeds the norm by a degree that is *significant* given our interests. 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 iff the other is. In other words, whenever you look at two objects in a Sorites sequence to see if you could find the boundary, you raise their similarity enough so that you could never find the boundary there; that's why we have a sense that the second premise of the paradox is true.

In the context of the current discussion, Graff's analysis would be implemented by characterizing the restrictions on the standard for (12a-b) as in (i).

(i) a. *λd.d is significantly greater than a norm of age for the comparison class* DOGS
b. *λd.d is significantly greater than a norm of age for the comparison class* PUGS

For the purposes of this paper, I will continue to assume the simple ordering in (13), but it may indeed be the case that Graff's analysis is ultimately the one we want to adopt. Most of what I say is independent of this decision, though the differences between relative and absolute adjectives, discussed in section 4 are perhaps more extreme if we adopt Graff's semantics for relative gradable adjectives.

[5]The comparison class should probably identify possible things of the relevant kind, rather than only the things that actually are part of the class in the context of utterance, but I will abstract away from this issue here.

(15a-c) require Spartacus to be a pug, as shown by the infelicity of (16a-c).[6]

(15) a. Spartacus is old for a pug.
b. Spartacus is not old for a pug.
c. Is Spartacus old for a pug?

(16) a. ??Barney the beagle is old for a pug.
b. ??Barney the beagle is not old for a pug.
c. ??Is Barney the beagle old for a pug?

We can capture these facts and implement the hypothesis that the crucial contextual parameter in gradable predicates in the positive form is the comparison class in an elegant way by adopting the following assumptions. First, let us assume that *for*-PPs that identify comparison classes are adjectival modifiers whose function is to restrict the domain of a GA to just those objects that are members of the set picked out by the comparison class. Such *for*-PPs will therefore have denotations along the lines of (17).

(17) $[\![\text{for a NP}]\!] = \lambda G_{\langle d,et\rangle}\lambda d\lambda x : x \in [\![\text{NP}]\!].G(d)(x)$

Let us further assume that even in the absence of an overt *for*-PP, the domain of the adjective can always be restricted in this way (e.g., by a null modifier whose denotation is a function of the denotation of the subject, as in Ludlow 1989). We may then revise our semantics for *pos* as in (18a), which I will abbreviate as in (18b).

(18) a. $[\![pos]\!] = \lambda G_{\langle d,et\rangle}\lambda x\exists d[d \succ mean\{d' \mid \exists y[y \in Dom(G) \wedge G(d')(y)]\} \wedge G(d)(x)]$
b. $[\![pos]\!] = \lambda G_{\langle d,et\rangle}\lambda x\exists d[d \succ norm(G) \wedge G(d)(x)]$

In short, the standard of comparison in the positive form of a GA is the mean degree of the relevant property for the objects in the (explicitly or implicitly restricted) domain of the adjective. On this view, the denotation of *pos* is itself fixed; contextual variability in the interpretation of gradable adjectives in the positive form comes from the possibility of shifting the domain of the adjective in different contexts, according to different comparison

[6]It should be observed that (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. Barney the beagle would be old for a pug.
b. If Barney the beagle were a pug, then he would be old for a pug.

classes. In other words, the approach outlined here maintains that contextual variability in the interpretation of vague predicates is just another example of contextual domain restriction, a phenomenon that is known to be pervasive in language (cf. the analyses of quantificational and modal domain restrictions in von Fintel 1994; Stanley 2002, and elsewhere). This approach therefore has the positive result of capturing the semantics of vague predicates in terms of a broader theory of contextual domain restriction.

Clearly, more details of this analysis need to be worked out, particularly with respect to examples involving implicit comparison classes, as well as its broader implications. The question I want to focus on here, though, is whether the general approach could possibly provide a general account of the full range of facts involving the interpretation of GAs in the positive form. Doing this requires answering two questions:

1. Can the contextual variability of the positive form always be explained in terms of domain restriction, and does the domain always correspond to a comparison class?

2. Do the truth conditions of the positive form always involve an ordering with respect to a norm in the sense defined above?

In sections 3 and 4, I will discuss two sets of data that suggest that the answer to both questions is 'no', and in section 5 I will propose an alternative analysis that will incorporate elements of the proposals outlined here, but will extend to the additional facts introduced below as well. Before moving to this discussion, however, I want to rule out an initially appealing alternative analysis to the one presented here.

2.3 Is the attributive form basic?

The examples up to now have all involved predicative forms of gradable adjectives, but such expressions also have attributive uses, as shown in (19).

(19) a. This is an expensive mission.
b. Spartacus is an old pug.

There is a strong intuition that attributive adjectives must use the nominals they modify as the basis for determining their comparison classes (and thereby fixing their domains, if we think in terms of the analysis presented in the previous section), leading some researchers to assume that the attributive form is basic, and that a predicative construction like (11) is really an elided form

of (19b) (see e.g. Montague 1970). If this were correct, then the semantic type of a gradable predicate would be somewhat different from what I suggested above; specifically, the comparison class would have to be an actual argument of the adjective (phrase).

On closer examination, however, there is reason to believe that this is not the right way to go, mainly because this analysis does not appear to buy us anything. Contrary to initial intuitions, attributive adjectives do not actually have to be interpreted with respect to a comparison class determined by the head noun. This is illustrated by three facts.

First, comparative forms of attributive adjectives, just like comparative forms of predicative adjectives, do not invoke a comparison class. (20a) can only be interpreted as comparing the relative costs of the motorcycle Benny drives and the car he drives; an interpretation along the lines of (20b) is impossible (and the sentence itself is not obviously grammatical).

(20) a. Benny drives a more expensive motorcycle than a car.
b. ??The motorcycle that Benny drives is more expensive for a motorcycle than the car he drives is for a car.

This indicates that a comparative attributive adjective at least is not relativized to a comparison class.

Second, attributive+nominal constructions do not give rise to the presuppositions we saw for predicative+*for*-PP constructions; in the attributive construction, the subject is merely entailed to have the property expressed by the nominal:

(21) a. That building is (actually) a tall tree.
b. That building isn't a tall tree.
c. Is that thing a tall tree? No, it's a tall building.

(22) a. ??That building is (actually) tall for a tree. (OK only if the building is a tree...)
b. ??That building isn't tall for a tree.
c. ??Is that thing tall for a tree? No, but it's tall for a building.

Finally, non-comparative forms of attributive adjectives can cooccur with measure phrases (if the adjective can be associated with a measure phrase in the first place):

(23) a. Benny is a three foot tall boy.
b. I need a two inch long piece of wire.

In contrast, *for*-PPs and measure phrases — as well as comparative morphology — are in complementary distribution.

(24) a. ??Benny is three feet tall for a boy.
b. ??Benny is taller for a three year old than he is for a Dutch child.
c. ??Benny is too tall for a three year old to wear these clothes.

If attributive adjectives necessarily restricted their domains to a comparison class identified by their associated nominals, then we would expect the examples in (23a-b) to be just as unacceptable as those in (24a-c).[7]

The conclusion to be drawn from these facts is that even in attributive constructions, we have to allow for the possibility that the standard of comparison is fixed independently of whatever comparison class the noun identifies. But this means that treating the attributive form as basic does not derive the comparison class-dependent interpretation for predicative uses; we still need something like the semantics of *pos* to do this for us even in the attributive form. We could assume that a default rule that restricts the domain of an attributive GA to that identified by the modified noun in order to account for our intuitions about the meanings of attributive adjectives, but we do not want to characterize the meanings of adjectives in general as attributive.

3 Varieties of standards of comparison

The goal of the next two sections is to show that the empirical domain that a comprehensive grammar of vagueness in gradable predicates needs to cover goes beyond the ability of the framework developed in section 2.2. I begin by considering examples which do not appear to involve a comparison class in the normal sense, although it may be the case that they can be treated under a broader notion of comparison class. In section 4, I turn to the case of 'absolute' gradable adjectives — adjectives whose standards appear to default to endpoints of a scale — which I argue cannot be adequately handled in terms of the semantic analysis presented in section 2.2.

[7] This argument may actually not go through. Assuming the analysis of *for*-PPs described in the previous section, the unacceptability of (24a-c) can be explained as follows. Unlike the *pos* morpheme, the various degree morphemes in (24) do not make use of the domain of the adjective as part of their meanings — they simple express orderings between degrees on a scale. Restricting the domain of the adjective therefore has no semantic effect, so the contributions of the *for*-PPs in (24) is completely vacuous.

However, even if we were to assume that the same thing is going on in (23a-b), it could be argued that these examples are acceptable because the nominals still make a semantic contribution: they provide the domains of the nominal predicates.

3.1 Standards of expectation

In her account of the Sorites Paradox, Graff (2000) argues for a semantic analysis of vague predicates much like the one outlined in section 2.2 (see note 4), in that it assumes that such predicates involve relations between objects and implicit standards. Graff also makes the important observation that there is more than one kind of standard. As an example, she points to different interpretations of predicates headed by the adjective *old*. Consider two dogs: Fido, who is fourteen, and Rover, who is twenty. In this context, the sentences in (25) could be used to express very different propositions.

(25) a. Fido is old.
b. Rover is old.

(25a) is most naturally interpreted to mean that Fido's age exceeds a mean of age for dogs, and so is interpreted exactly as expected according to the principles laid out in section 2.2. (25b) can be understood to mean something stronger, however: Rover's age exceeds a 'standard of life expectancy' for dogs. Using the former standard, both sentences in (25) are true; using the latter standard, (25a) is false.

A second case discussed by Graff involves standards based on a 'standard of expectation'. ((25b) may be a variant of this sort of case.) Consider a context in which I have a friend whose child is definitely not tall for his age, is in fact fairly small. If I visit my friend after not seeing her for some time and discover that her child has grown more than I expected, but is nevertheless still fairly small, it is perfectly felicitous for me to assert (26).

(26) Benny is tall!

Here the interpretation is not 'Benny's height exceeds the mean of height for children like him', but rather something like 'Benny's height is greater than I thought it would be'. Exclamative constructions like 'How tall you are!' represent a grammaticization of this sort of interpretation.

The problem that these examples present is that they appear to require us to use something other than 'norm for a comparison class' to compute the truth conditions of these examples. If this is right, then it calls into question both the hypothesis that the standard is a 'norm' (in the sense defined above) and the idea that the crucial parameter of contextual variation in a gradable predicate is its comparison class-restricted domain.

There is a potential way out of this problem, however, which is to assume that examples like (26) and (25b) involve somewhat more arbitrary

comparison classes than those we have considered so far, such as the ones indicated in (27a-b).

(27) a. Rover is old **for an old dog**.
b. Benny is tall **for a boy of the height I was expecting**.

The question that this hypothesis has to answer is whether there are any constraints on the composition of comparison classes, or whether in fact any arbitrary property will do, if it can be recovered or made salient in the context.

3.2 Standards of differentiation

Another potential problem for the idea that variability is always a function of a comparison class comes from examples illustrating what I will call 'standards of differentiation'. This sort of case is recently discussed by Kyburg and Morreau (2000) (see also Sedivy, Tanenhaus, Chambers, and Carlson 1999), who show that gradable adjectives can be used to distinguish one object from another, even when the degree to which that object possess the relevant property is less than the mean for the appropriate comparison class in the context of utterance.

Consider, for example, a context in which one farmer is negotiating with another farmer over two pigs (adapting Kyburg and Morreau's example). One of the pigs is a runt, the other is bigger, but neither truly qualifies as big (for a pig). It is nevertheless the case that a definite description like *the big one* can be quite naturally used to identify the larger of the two pigs, since the standard for *big* can be shifted just for the purpose of differentiating one pig from another. Thus (28a) could be both felicitous and true in this context, even if (28b) is false.

(28) a. The big one can talk to spiders.
b. The one that can talk to spiders is big for a pig.

The problem is that (28a) clearly does not make reference to a norm of bigness for the comparison class PIGS. Again, however, this sort of example can be subsumed under the general approach outlined in section 2.2 if we assume that the domain of an adjective can be restricted to an arbitrarily small set of objects, in this case the two pigs under discussion. Indeed, this is an option for quantifiers: in the context described above, *each* in (29) could be construed as quantifying just over the two pigs in the pen, rather than a larger set.

(29) Each pig received its vaccinations.

3.3 Interest-relative standards

A third potential problem for the general approach outlined in section 2.2 comes from examples in which the comparison class appears to be clearly indicated, but speakers may still legitimately disagree on the truth conditions contributed by a GA in the positive form. As an illustration, imagine a context in which two individuals are buying coffee at Starbuck's. They pay exactly the same amount for their coffees, and moreover they have the same past experience and acquaintance with coffee and coffee drinking. However, these two individuals may still disagree on the truth of (30) if, for example, one is poor and thrifty and the other is rich and extravagant.

(30) A cup of coffee at Starbuck's is expensive (for urban coffee joints).

What seems to matter here is not what the average price of coffee in urban coffee joints is, but rather what the interests of the participants are *vis a vis* the amount of money they are willing to spend on coffee. In other words, it seems most reasonable to assume here that the two coffee drinkers make the same assumptions about the comparison class, and so fix the domain of the adjective in the same way, but that they end up with different standards of comparison — and different truth conditions — because they use different kinds of 'norm' functions.

This kind of example could be accounted for by modifying our semantics of *pos* along the lines suggested in Graff 2000, so that the standard of comparison is required to be not merely greater than the mean for the domain of the adjective, but rather significantly greater than the mean, where what counts as significant is a function of the interests of ... (the speaker??) The question we then need to answer is whether Graff's 'significantly greater than' relation is operative in all predications involving GAs. If so, then we can maintain the proposals presented in the previous section, modified as suggested here, as a general semantic analysis for the positive form. If we find that this relation is not appropriate for some cases, however, then we will need to explain how it comes to be part of the meaning of examples like (30).

3.4 Summary

Further research will undoubtably uncover still more varieties of standards; whether they can all be subsumed under a version of the analysis of the positive form presented in section 2.2, possibly modified along the lines suggested in Graff 2000, remains to be seen. At the moment, we should take the case

studies presented above as potential indications that the analysis outlined in section 2.2 may not be fully general, but not as clear counterexamples to it. In the next section, I turn to more clearly problematic cases.

4 Absolute gradable adjectives

Most of the literature on vagueness in gradable predicates assumes (implicitly or explicitly) that all gradable predicates give rise to the sort of contextual variability we have seen with *old*, *big*, *expensive* and so forth, but this is actually not the case. In addition to the class of RELATIVE gradable adjectives — GAs of the sort we have been considering so far, whose interpretations in the positive form appear to involve reference to a standard of comparison computed as a mean over a comparison class-restricted domain, as outlined above — there is a class of adjectives that are demonstrably gradable but whose interpretations are not context-dependent in the way we have seen so far.

For example, the adjectives in (31) simply require their arguments to possess some minimal degree of the property they describe, not that the degree to which the arguments possess this property is greater than some contextually determined standard of comparison.

(31) *Minimum standard absolute adjectives*

a. The baby is awake.
b. The table is wet.
c. The door is open.
d. The rod is bent.

Under normal usage, (31a) does not mean that the degree to which the baby is awake surpasses some standard (for babies), but rather simply means that the baby has a non-zero level of awakeness. Likewise, (31b) is true as long as there is some amount of water on the table, (31c) just requires some minimal positive aperture of the door, and (31d) is true of a rod that is minimally bent.

The adjectives in (32) are similar, except that their arguments are required to posses a maximal degree of the property in question.

(32) *Maximum standard absolute adjectives*

a. The glass is full.
b. The road is flat.
c. The door is closed.
d. The rod is straight.

(32a) typically means that the glass is completely full, not that its contents

fall above some standard of fullness, (32b) is an assertion that the road has no bumps, (32c) requires the door to be completely closed, and (32d) requires a completely straight rod.

Following Unger (1975) and more recently Kennedy and McNally (1999a, 2004), I will refer to adjectives like those in (31) and (32) as ABSOLUTE (gradable) adjectives (see also Yoon 1996 and Rotstein and Winter to appear). Absolute adjectives are gradable, as shown by their acceptability in comparatives:

(33) a. The baby is more awake now than it was a few minutes ago.
b. The table is wetter than the floor.
c. The door is more open than it should be.
d. Rod A is more bent than Rod B.

(34) a. My glass is fuller than your glass.
b. This road is flatter than that one.
c. The door should be more closed than it is.
d. Rod B is straighter than Rod A.

If both relative and absolute gradable adjectives are of the same semantic type — expressions of type $\langle d, et\rangle$, which their similar behavior and interpretations in comparative constructions suggests — then the truth conditions associated with the absolute adjectives in (31) and (32) are unexpected. According to the theory presented in section 2.2, the positive form of a gradable adjective necessarily involves reference to a comparison class, since the only way to saturate the degree argument of the adjective in the positive form is to use the null morpheme *pos* which, by hypothesis, contains the comparison class variable as part of its meaning. The examples in (31) and (32) do not involve interpretations with respect to comparison classes, however, and so counterexemplify this prediction and call the broader semantic analysis into question.

This problem presented by absolute adjectives has not been explicitly addressed in previous work on the semantics of gradable predicates, presumably because there is a strong initial intuition that the adjectives in (31) actually require something significantly more than a minimum standard, and that those in (32) actually allow something less than a maximum standard. These intuitions are supported by examples like those in (35).

(35) a. I'm not awake yet.
b. The tank is full, but you can still top it off. It's not completely full yet.

c. The theatre is empty tonight.

(35a) can be felicitously uttered by someone who is not talking in his sleep. Likewise, most speakers I have consulted feel that *full* only requires its arguments to fall *near* the maximal value on the scale, pointing to examples like (35b), which does not sound contradictory (but cf. the examples discussed below in (37)-(38)). Similarly, (35c) can be used to describe a situation in which only a very few people show up to a film in a very large movie theatre.

On the whole, it is fairly easy to come up with other 'imprecise' uses of absolute limit adjectives. However, as argued in detail by Kennedy and McNally (2004), there are both coherent theoretical reasons and compelling empirical arguments for the conclusion that absolute adjectives, unlike relative adjectives, do in fact involve maximum and minimum standards, rather than standards based on the mean for some comparison class.

One possible theoretical explanation for imprecise uses of absolute adjectives would be to maintain that the propositions conveyed by the sentences in (35) are strictly speaking false in the contexts described, and explain their felicity and informativity in terms of general pragmatic principles governing the interpretation of 'loose talk'. (This is essentially Unger's position.) This idea can be implemented in terms of Lasersohn's (1999) theory of PRAGMATIC HALOS, which provides a framework for determining how much deviation from what is actually true counts as 'close enough to the truth' in any context.

A second possibility would be to claim that precise vs. imprecise uses involve different *granularities* of degrees, and that the same object may count as maximally full at coarse granularities but not a finer granularities. The relevant analogy is to a meter stick that indicates only centimeters vs. one that also indicates millimeters. If we assume that an object's length is at least as great as degree n just in case it falls in the space between n and $n-1$, then an object that actually measures 995 milliliters will count as maximal with respect to the first stick but not with respect to the second.

4.1 Empirical evidence for the absolute/relative distinction

Regardless of the particular theoretical account of imprecise uses of absolute adjectives, what is crucial is that there are several important empirical distinctions between absolute and relative adjectives, which all point towards the latter having endpoint-oriented standards of comparison. Here I focus on five crucial differences.

4.1.1 Entailments

The first piece of evidence that distinguishes absolute adjectives from relative ones comes from entailment patterns in the positive and comparative forms. If the standards associated with the former involve maximum and minimum values, then the denotations of the predicates they head can be characterized informally as in (36a) (for minimum standard adjectives) and (36b) (for maximum standard adjectives).

(36) a. $[\![[_{\text{AP}}\ \text{A}_{min}]]\!] = \lambda x.x$ has more than a minimal degree of A-ness
b. $[\![[_{\text{AP}}\ \text{A}_{max}]]\!] = \lambda x.x$ has a maximal degree of A-ness

If this is correct, we should find the following entailments patterns. First, (36a) predicts that a negative assertion of the form *x is not* A_{min} should entail that x possesses no amount *adj*-ness at all (assuming that the minimal degree on a closed scale represents a zero amount of the relevant property). The contradictory statements in (37) illustrate that this prediction is borne out.

(37) a. #My hands are not wet, but there is some moisture on them.
b. #The door isn't open, but it is ajar.

Second, (37b) 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 to the adjective. When we do this, as in (38), we see that the expected entailments arise:

(38) 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 entail only that its argument falls above whatever standard of comparison is calculated based on the context. As a result, neither of the entailments discussed above should hold. This is correct:

(39) 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 to appear). As shown by the examples in (40), there exist pairs of antonyms such that negation of one form entails the assertion of

the other:

(40) 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 (40) 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 negative degree on the same scale (see Kennedy 2001), the entailment relations in (40) follow from the truth conditions in (36).

Relative antonyms do not show the same entailment relations:

(41) a. The door is not large. ⇏ The door is small.
b. The table is not expensive. ⇏ The table is inexpensive.
c. The baby is not energetic. ⇏ 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 crucially need not be endpoints (in fact, cannot be endpoints if the scales are open). Since a context dependent standard is determined for particular uses of particular adjectives, it need not be the case that the standard for e.g. *large* be the same as that of its antonym *small*, and we allow 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'). Indeed, the possibility of such 'borderline cases' is one of the defining properties of vague predicates; see Williamson 1994 for general discussion.

Turning to comparative constructions, if the truth conditions associated with minimum and maximum standard absolute adjectives are as in (36), we expect them to generate positive and negative entailments to the positive form, respectively (depending on which argument of the comparative we use). This prediction is borne out, as shown by the examples in (42)-(43).

(42) a. The floor is wetter than the countertop. ⇒
b. The floor is wet.

(43) a. The floor is drier than the countertop. ⇒
b. The countertop is not dry.

Assuming the comparative imposes an asymmetric ordering on its arguments (see the truth conditions for the comparative in (5)), (42a) will be true only

if the floor has some degree of wetness (though the countertop may have a zero degree). The truth of (42b) then follows directly from (36a). Similarly, in order for (43a) 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 (36b), 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 (44)-(45).

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

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

This also follows, since 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 determined standard of comparison.

4.1.2 For-PPs

A second difference between relative and absolute adjectives involves the distribution of *for*-PPs. We have already seen that such expressions can be used to introduce the comparison class with respect to which a context-dependent standard is determined:

(46) 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.

This type of *for*-PP is often infelicitous with absolute adjectives, however. This follows if the interpretation of these adjectives typically does not involve reference to a standard computed on the basis of restricted domain for the adjective: the for-PPs in (47) contribute nothing to the assertion.[8]

[8]Graff (2002) suggests that the problem with the examples in (47) is that the function that returns a mean is not defined for the comparison classes in these examples. This seems to restate the problem, however; we still have the question of why e.g. *a piece of outdoor furniture* can be used to find a mean for *small* but not for *wet*. I suspect that if we really push the context, we can make even the examples in (47) felicitous, but we also want to

(47) 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.
c. ??My glass is {full, empty} for a wine glass.
d. ??The door is {closed, open} for an office door.

This restriction is not absolute, however. The examples in (48) involve absolute adjectives, but the *for*-PPs are felicitous. (Thanks to Jeff King for bringing these to my attention.)

(48) 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 a comparison class to shift the standard of a maximum standard absolute adjective. (It seems to be a lot more difficult to construct similar examples for minimum standard adjectives, a point I return to below.) However, I do not think that these examples call into question the basic relative/absolute distinction; rather, they can be used to reinforce the initial claim that such a distinction should be made. An important difference between the examples in (48) and those in (46) is that (48a-b) strongly implicate the negations in (49): (48b), for example, would be infelicitous if used to describe a theater that is in fact completely empty.

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

This is not generally true of relative adjectives with *for*-PPs, as illustrated by the examples in (50).

(50) a. Nigel is tall for a dwarf.
b. Nigel is tall for a basketball player.

Given what we know about dwarfs, we would likely infer from (50a) that Nigel is not tall. In contrast, given what we know about basketball players, we would naturally infer from (50b) that Nigel is tall. The examples in (50) show that whether or not we make an inference from '*x is A for a y*' to '*x is (not) A*', when *A* is a relative adjective, is dependent on the nature of the

know what is responsible for the base-level contrast between the examples in (47) and those in (46). The proposal I develop in section 5 will answer this question: we *need* a restricted domain to get an informative proposition with relative adjectives, but we don't need one to get an informative proposition for absolute adjectives.

comparison class identified. In principle, the standard determined on the basis of this comparison class could fall either above or below whatever standard is determine based on properties of the subject alone, allowing for either a positive or negative inference to the simple sentence without the comparison class, respectively. In other words, with relative adjectives, a *for*-PP can move a standard up or down, relative to whatever standard would be identified in the absence of such information.

The fact that examples like (48) always give rise to the negative inferences in (49), on the other hand, indicates that the comparison class can only move the standard below what it would have been without the comparison class. If the 'default' standard for absolute adjectives like *straight, full* etc. is the maximal value on the relevant scale, then the only alternative values would in fact be lower ones, and the facts observed here are exactly what we would expect to see.

Extending this line of reasoning a bit further, we can make the following generalization: if the default standard for absolute adjectives like *straight* and *full* is a maximum value on the relevant scale, then a *for*-PP whose function is to explicitly identify a comparison class should be felicitous only if it has the effect of moving the standard away from the maximum (i.e., down the scale). If, on the other hand, such absolute adjectives do not default to a maximum standard — if they are just a special case of relative adjectives, whose standards happen to tend to fall near the upper end of a scale — then we should also be able to use a *for*-PP to explicitly indicate that the standard is close to or at the maximum. The following minimal pair in provides a test of these two hypotheses:

(51) a. That groove is straight for a groove carved by such an old machine.
b. ??That groove is straight (even) for a groove carved by such a precise machine.

The fact that (51b) is infelicitous suggests that it is not possible to use a comparison class to get closer to the maximum value on the scale, indicating that the maximum is in fact the default.

4.1.3 No standard of differentiation

The preceding array of facts indicated that it is in general difficult to shift the standard of an absolute adjective, and even when this is possible, the standard can only be shifted in one direction, which follows if the default is an endpoint-oriented standard. Another illustration of this restriction comes from

the use of gradable adjectives in definite descriptions. As discussed above (see (28)), relative adjectives can have 'standards of differentiation' in such contexts. Absolute adjectives do not appear to permit this sort of interpretation, however.

Consider a context in which two glasses of beer are on the table, one of which is half full and one of which is 2/3 full. Referring to the latter with the definite description *the full glass*, as in (52a), is infelicitous; instead, it is necessary to use the comparative form of the adjective as in (52b). (The comparative form is of course also possible with relative adjectives when a distinction is being made between two objects, but it is not required.)

(52) a. #The full glass of beer is mine.
b. The fuller (of the two) glass(es) of beer is mine.

Minimum standard absolute adjectives behave the same. If A and B are standing in front of two partially open doors, one that is barely open and one that most of the way open, A cannot felicitously direct B towards the more open of the two doors by saying (53a); A must say (53b).

(53) a. #You should go through the open door.
b. You should go through the more open (of the two) door(s).

These facts follow if the standards for *full* and *open* are fixed at the maximum and minimum values of the respective scales (modulo imprecision). Since the standards cannot be shifted, the existence and uniqueness presuppositions associated with the definite descriptions in these examples (that there is a full glass of beer/open door) are not satisfied, and the (a) sentences are anomalous.

4.1.4 Pretty

The degree modifier *pretty* also distinguishes between relative and absolute gradable adjectives. When it modifies a relative adjective, it can have a meaning similar to *very*: it 'boosts' the value of the standard of comparison. On this reading, (54a) thus entails (54b).

(54) a. The rod is pretty long.
b. The rod is long.

Pretty can also have a 'hedging' interpretation, so that (54a) means something like *the length of the rod is close to the standard* — maybe above, maybe below; we can't be sure. On this reading, (54a) doesn't commit the speaker to (54b).

With absolute adjectives, the story is a bit different. Minimum standard absolute adjectives seem to allow both interpretations of *pretty*, but the 'hedge' use still entails the positive form:

(55) a. The rod is pretty bent.
b. The rod is bent.

This follows if the standard of comparison is a minimal value on the scale: the only direction of deviation from the standard is up, so even a little bit of an upwards hedge will trigger an entailment to the positive form.

When *pretty* modifies a maximum standard absolute adjective, only the hedging use is possible, and we get an entailment to the negation of the positive form (cf. Unger 1975):

(56) a. The rod is pretty straight.
b. The rod is not straight.

This follows if the standard of comparison is a maximum value on the scale: the standard-boosting interpretation should be ruled out, and only the 'close to but not quite there' understanding of the hedge version should be possible. Again, if the standard of comparison for absolute adjectives could be arbitrarily fixed relative to a comparison class, we would have no explanation for this systematic entailment pattern.

4.1.5 The Sorites Paradox

A final difference between relative and absolute adjectives concerns their interaction with the Sorites Paradox. Typical examples of the paradox are arguments like the following, which uses the relative adjective *big*.

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

The problem with this argument is that it should be valid, but the conclusion is clearly unjustified (see Graff 2000 for comprehensive discussion and a potential solution to the problem). One type of explanation for the paradox, is that even though there is some cutoff point between the big theaters and the non-big ones, which would render the second premise invalid, is that we can never actually know where this point is. Thus we are willing to accept the second premise even though we shouldn't, which leads to the problem of the false conclusion.

In the case of absolute adjectives, however, we do know where the cutoff point is: at the minimum or maximum end of the scale, depending on the adjective. If this is right, then the prediction is that absolute adjectives should not give rise to the paradox, and this appears to be correct: (58) does not lead to a paradoxical conclusion, because already we may object to the truth of the second premise.

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

4.2 Summary

The conclusion to be drawn from the facts discussed here 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 mean for a comparison class, and in the majority of cases, it clearly isn't. But this in turn means that a semantics for gradable predicates that always analyzes the standard of comparison in the positive form in terms of a mean for a comparison class-restricted domain fails to generate the correct truth conditions in all cases. More generally, it means that a 'grammar of vagueness' that attempts to handle the contextual variability of gradable predicates strictly in terms of the identification of a comparison class is empirically inadequate.[9]

5 Towards a grammar of vagueness

The question that I started with is the following: how are sentences of the form *x is GrAdj* assigned determinate truth conditions in a context of utterance? What we have seen is that there are at least three semantically distinct sorts of interpretations that can be assigned to a predicate headed by a gradable adjective, depending on whether the adjective is relative, a minimum standard absolute, or a maximum standard absolute; these are summarized in (59).

(59) $[\![[_{\text{AP}}\ \text{A}]]\!] =$

[9] At the same time, if Graff's (2000) solution to the Sorites Paradox is correct (see note 4), then her hypothesis that an object must possess a degree of property A that is 'significantly greater' than the standard can only work for relative adjectives: this would arguably give us truth conditions that are too strong for minimum standard absolute adjectives, and it would clearly give us contradictory truth conditions for maximum standard absolute adjectives.

a. $\lambda x.\exists d[d \succ norm(Dom(A)) \wedge [\![A]\!](d)(x)]$ *relative A*
b. $\lambda x.\exists d[d \succ min(\text{SCALE}(\text{A})) \wedge [\![A]\!](x)(d)]$ *minimum absolute A*
c. $\lambda x.\exists d[d = max(\text{SCALE}(\text{A})) \wedge [\![A]\!](x)(d)]$ *maximum absolute A*

Furthermore, if it turns out that the various standards of comparison discussed in section 3 cannot be accounted for in terms of something like (59a), we will need further denotations for relative adjectives.

There are at least two ways to go about explaining how we get these different sorts of interpretations. The first is to simply assume that the two classes have different semantic analyses, i.e., are of different semantic types. Specifically, we could assume that relative adjectives have the semantic analysis outlined in section 2.2: they are expressions of type $\langle d, et\rangle$ and must combine with the *pos* morpheme in the positive form to generate a property of individuals, with the result that their interpretation requires fixing a comparison class. Absolute adjectives, in contrast, do not combine with *pos*, but instead somehow directly encode the appropriate relation to the standard of comparison (minimum or maximum) in their lexical entries. This amounts to saying that absolute adjectives are not gradable in the technical sense — they don't have a degree argument that needs to be saturated — but rather can directly encode properties of individuals.

This claim is clearly false in a strong sense, since absolute adjectives combine with comparative morphology, unlike nongradable adjectives like *extinct*:

(60) a. The table is {flatter, wetter} than the chair.
b. The bathtub is {fuller, emptier} than the sink.
c. ??The wooly mammoth is more extinct than the elephant.

Assuming that acceptability of the comparative form indicates that an expression has a degree argument (see Klein 1991; Kennedy 1999b), these facts force us to conclude that absolute adjectives are — or at least can be — of the same semantic type as relative adjectives. This means that our first option for explaining the semantic differences between relative and absolute adjectives amounts to saying that the former are always of type $\langle d, et\rangle$, while the latter are ambiguous between type $\langle d, et\rangle$ meanings (which show up in comparatives) and type $\langle et\rangle$ meanings (the positive form).

The second option for explaining the relative/absolute distinction maintains the assumption that all gradable adjectives — relative and absolute alike — have the same semantic type and same basic meaning, namely the one stated in (61): they describe relations between individuals and degrees on a scale defined by the adjective; see the discussion of (3) in section 2.1).

(61) $[\![A]\!] = \lambda d \lambda x.\mathbf{m}_A(x) \succeq d$

However, this approach requires us to revise our analysis of the positive form so that it derives the appropriate standard of comparison for the appropriate adjectival input based on some feature or features of the base adjective.

There are a couple of compelling reasons for taking this strategy, in addition to the unexplanatory nature of the 'ambiguity analysis' outlined above. First, as made clear in particular by the discussion in section 4.1.2, even absolute adjectives show variation in their standards, albeit a much more limited range of variation than relative adjectives (the standard can only be moved in one direction, and typically not very far). In other words, relative and absolute adjectives are siblings, with fundamentally the same type of meanings, not distant cousins.

Second, when we take a broader look at the semantic properties of relative and absolute gradable adjectives, we see that there is in fact a semantic feature that distinguishes the two classes: the structure of the scales onto which they map their arguments. Kennedy and McNally (2004) show that one parameter of variation among gradable adjectives is whether they map their arguments onto an open scale (no maximal/minimal values) or a closed one (maximal and/or minimal values) (see also Kennedy and McNally 1999b; Paradis 2001; Rotstein and Winter to appear). Variation in scale structure corresponds to variation in the range of the measure function. In particular, Kennedy and McNally argue that at least the four scale types schematically represented in (62) are attested.

(62) *A typology of scale strucures*

a.	$\mathrm{A}_{(0,\infty)}$	*(totally) open*
b.	$\mathrm{A}_{[0,\infty)}$	*lower closed*
c.	$\mathrm{A}_{(0,1]}$	*upper closed*
d.	$\mathrm{A}_{[0,1]}$	*(totally) closed*

Empirical evidence for at least this four-way typology of scale structures is based on the acceptability of modifiers that pick out maximal degrees and the assumption that the maximal value for a positive adjective is the 'upper' end of the scale, while that of a negative adjective is the 'lower' end of the scale (Kennedy 2001). Given this, we predict the following pattern of modifier+adjective acceptability:

(63)

	OPEN	L-CLOSED	U-CLOSED	CLOSED
$[\mathrm{Deg}_{max}\ \mathrm{A}_{pos}]$	??	??	√	√
$[\mathrm{Deg}_{max}\ \mathrm{A}_{neg}]$	??	√	??	√

The maximizing modifier *absolutely* provides particularly clear judgments in this test (cf. Unger 1975; *almost* shows a similar pattern, as shown in Rotstein and Winter to appear). As shown by the following examples, the expected pattern does in fact emerge:

(64) *Open scales*

a. ??absolutely {tall, deep, expensive, likely}
b. ??absolutely {short, shallow, inexpensive, unlikely}

(65) *Lower closed scales*

a. ??absolutely {possible, bent, bumpy, wet}
b. absolutely {impossible, straight, flat, dry}

(66) *Upper closed scales*

a. absolutely {certain, safe, pure, accurate}
b. ??absolutely {uncertain, dangerous, impure, inaccurate}

(67) *Closed scales*

a. absolutely {full, open, necessary}
b. absolutely {empty, closed, unnecessary}

The facts here are somewhat counterintuitive: one might think that the scale used by the antonym pair *long/short*, for example, is closed on the bottom (i.e., has a 'zero length' value). The infelicity of expressions like *??absolutely short/young/inexpensive* and so forth, however, shows that semantically there is no such minimal value. Despite initial intuitions, then, the linguistic data indicate that those in (67) use scales without minimal or maximal values.

What is important to the current discussion is Kennedy and McNally's observation that there is a systematic relation between scale structure and standard of comparison: as seen by the examples in (64)-(67), gradable adjectives associated with totally open scales are relative; gradable adjectives that use totally or partially closed scales are absolute. The first generalization is exceptionless: since open scales lack endpoints, it is impossible for open scale adjectives to have endpoint-oriented standards. The second generalization is not exceptionless (as we have seen), but it does appear that the standards for closed-scale adjectives default to an endpoint of the scale: the minimum in some cases (e.g., *wet* and *open*); the maximum in others (e.g., *dry* and *straight*).

Crucially, there isn't a logically necessary reason why adjectives with closed scales should have absolute standards. The fact that we do get such a strong tendency in this direction indicates that this feature gives rise to (potentially overridable) constraints on the value of the standard of comparison.

The question that we need to answer now is why we end up with maximum or minimum standards in these cases, and why we end up with context-dependent standards in the case of relative adjectives.

The answer that I would like to consider here is that the basic meaning of the positive form is not a relation to a 'norm', as I suggested in section 2.2, but rather just involves restricted existential quantification over degrees, as stated in (68).

(68) $[\![[_{Deg}\ pos]]\!] = \lambda G_{\langle d,et\rangle}.\lambda x.\exists d[\mathbf{R}(d) \wedge G(d)(x)]$

The empirical observations we have made so far indicate the following pattern for fixing the value of the restriction variable $\mathbf{R}$:[10]

(69) a. If A uses a lower closed scale, $\mathbf{R} = \lambda d.d \succ min(\textsc{scale}(A))$
b. If A uses an upper closed scale, $\mathbf{R} = \lambda d.d = max(\textsc{scale}(A))$
c. If A uses an open scale, then, $\mathbf{R} = \lambda d.d \succ norm(dom(A))$

The final piece of the puzzle is ensuring that we get the right value of $\mathbf{R}$ for the right adjective. This is not a problem for open scale adjectives: since the interpretations in (69a-b) are defined only for scales with endpoints, these will not be options in the first place. In effect, only the 'domain dependent' interpretation is available for open scale adjectives, so the prediction is that these will necessarily have context-dependent interpretations.

The problem is ensuring that (69c) is not an option — or at least a marked one — for closed scale adjectives. But if the proposed semantics for *pos* in (68) is correct — if the positive form includes an implicit restriction variable — then what rules out interpretations like (69c) for closed scale adjectives? In other words, what *forces* absolute/closed-scale adjectives to have maximum/minimum standards, even just as defaults?

I would like to suggest that this follows from a general principle of 'interpretive economy', which regulates the way that the values of free variables are determined. In short: if it is possible to fix the value of a variable strictly

[10] (ic) represents a 'norm for a comparison class', in the sense discussed in section 2.2. However, to account for the different types of standards observed in section 3 for relative adjectives, we may want to allow for further options for the value of $\mathbf{R}$ when A is an open scale adjective, such as (at least) the following:

(i) a. $\lambda d.d \succ \iota d : [\![A]\!](x)(d)$ (for some salient x; *norm of differentiation*)
b. $\lambda d.d \succ d'$ (for some salient d'; *arbitrary standard*)
c. $\lambda d.d! \succ norm(Dom(A))$ ($! \succ$ = Graff's significantly greater than; *interest-relative norm*)

based on properties of the conventional meaning of expressions in a sentence (or logical form), then the value of the variable must be fixed in this way. Contextual assignment of values to variables is a 'last resort', available only when the conventional meaning of expressions in the sentence does not help.

This sort of principle needs to be qualified in a crucial way, however. For example, the simplest way to fix the value of the restriction variable **R** in the positive form of a gradable adjective would be to just use the mere fact that *pos* is a quantifier over degrees to set **R** to some vacuous property, e.g. *is a degree on the relevant scale*. The resulting truth conditions for a sentence of the form *x is A*, would then be equivalent to *x has some degree of A-ness*. However, in order to account for basic selectional restrictions of gradable adjectives, we need to assume that only objects that have some degree of the relevant property are in the domain of the adjective in the first place, which means that this way of fixing the value of **R** results in a proposition that could not fail to be true.

What we need, then, is an interpretive economy principle along the lines of (70).[11]

(70) *Avoid Context Dependence*
If a contingent proposition can be derived strictly on the basis of lexical properties of elements of the sentence, then do so.

In the case of the positive form, *Avoid Context Dependence* dictates that if it is possible to use the lexical properties of a gradable adjective to determine the standard of comparison — to fix the value of the degree restriction **R** — then this is what should be done. In the case of GAs with closed scales, the minimum or maximum values (or both) provide 'fixed' reference points with respect to which the standard can be computed, resulting in the values for **R** stated in (69a-b) above. The scalar properties of open scale adjectives provide no such reference points, however, so the only way to get informative truth conditions for the positive form is to restrict the degree quantifier based on some contextual information, such as the implicit domain of the adjective. Moreover, if *Avoid Context Dependence* is essentially a pragmatic constraint, and so is violable under certain conditions, then we expect relative-like interpretations of absolute adjectives to be in principle possible, though the absolute

[11] An obvious question raised by the principle in (70) is how it should apply to other cases of contextual domain restriction, in particular quantificational domain restriction. In particular, do the semantic properties of the nominal argument of a quantificational determiner ever affect the possible value of additional implicit domain restrictions on the determiner? This is clearly an issue that needs to be resolved.

interpretation should be the default. This is consistent with the facts we have observed.

Avoid Context Dependence makes the right predictions about GAs with open scales, upper closed scales, and lower closed scales, but what about adjectives which have totally closed scales and maximum standards? As things stand, the prediction is merely that these adjective should have should be absolute; the proposals do not determine whether the standard should be a minimum or maximum. Kennedy and McNally (2004) argue that in the case of deverbal adjectives with totally closed scales, the standard is computed as a function of a different lexical property of the adjective: role played by the argument of the adjective in the corresponding event. This proposal does not extend to lexical adjectives like *full, empty*, however, which quite clearly have maximum standards.

One difference between a minimum standard interpretation and a maximum standard one is that the latter entails the former, but not vice versa. This means that the maximum standard interpretation is stronger: its true in fewer situations. This observation suggests a possible preference for maximum standard interpretations of adjectives with totally closed scales, and the role of a second pragmatic constraint:

(71) *Maximize Strength*
If A and B are competing interpretations for a structure S, and A is stronger than B, then the interpretation of S is A.

However, it also appears that there is some variability in the interpretation of such adjectives. According to the diagnostics discussed above, *open* uses a totally closed scale (*absolutely open/closed*), but it appears to be able to have both minimum and maximum standard interpretations in different contexts. For example, if we are members of the crew of the starship *Enterprise*, we are most likely to understand (72a) in a minimum standard sense and (72b) in a maximum standard sense. (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.)

(72) a. The airlock is open and the cabin is depressurizing!
b. The space door is open, so we can take the ship out of the station.

6 Concluding remarks

The answer to the question we started out with, then, is the following: determinate truth conditions for the positive form of a gradable adjective are derived by existentially quantifying over the degree argument of the adjective;

like other types of quantification, this involves an implicit variable over the domain of quantification. The way that the value of this variable is determined is subject to a general constraint on interpretive economy: if its value can be determined strictly on the basis of lexical — scalar — properties of the adjective, then it must be so determined. Only when this is not possible (or when other factors override this option) can the restriction be fixed contextually; this option is a 'last resort'.

There are many significant questions that remain unanswered, but perhaps the most important of them is the following: if the standard of comparison is determined contextually for open scale adjectives, why does it so often correspond to a norm for a comparison class? Although the data discussed in section 3 suggested that other sorts of standards are possible, it is nevertheless the case that a comparison class-based standard — more or less along the lines of what I proposed in section 2.2 — is a kind of default for open scale/relative adjectives. One possibility is that there is some sort of cognitive answer to this question, e.g., that such standards reflect typical distributions of objects in arbitrary domains with respect to gradable properties. But I don't have a serious answer to this question here.

References

Barker, Chris. 2002. The dynamics of vagueness. *Linguistics and Philosophy* 25:1–36.

Bartsch, Renate, and Theo Vennemann. 1973. *Semantic structures: A study in the relation between syntax and semantics*. Frankfurt: Athäenum Verlag.

Bierwisch, Manfred. 1989. The semantics of gradation. In *Dimensional adjectives*, ed. Manfred Bierwisch and Ewald Lang, 71–261. Berlin: Springer-Verlag.

Cresswell, M. J. 1977. The semantics of degree. In *Montague grammar*, ed. Barbara Partee, 261–292. New York: Academic Press.

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

Fine, Kit. 1975. Vagueness, truth and logic. *Synthese* 30:265–300.

von Fintel, Kai. 1994. Restrictions on quantifier domains. Doctoral Dissertation, University of Massachusetts, Amherst.

Fintel, Kai von. 1994. Restrictions on quantifier domains. Doctoral Dissertation, University of Massachusetts at Amherst.

Gawron, Jean-Mark. 1995. Comparatives, superlatives and resolution. *Linguistics and Philosophy* 18:333–380.

Graff, Delia. 2000. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics* 20:45–81.

Graff, Delia. 2002. Comments on Christopher Kennedy's 'The landscape of vagueness'. Presentation at University of Michigan Workshop on Semantics and Philosophy of Language, November 8, 2002.

Heim, Irene. 1985. Notes on comparatives and related matters. Unpublished ms., University of Texas, Austin.

Heim, Irene. 2000. Degree operators and scope. In *Proceedings of SALT X*. Ithaca, NY: CLC Publications.

Hellan, Lars. 1981. *Towards an integrated analysis of comparatives*. Tübingen: Narr.

Kamp, J.A.W. 1975. Two theories of adjectives. In *Formal semantics of natural language*, ed. Edward Keenan, 123–155. Cambridge: Cambridge University Press.

Kennedy, Christopher. 1999a. *Projecting the adjective: The syntax and semantics of gradability and comparison*. New York: Garland. (1997 UCSC Ph.D thesis).

Kennedy, Christopher. 1999b. *Projecting the adjective: The syntax and semantics of gradability and comparison*. New York: Garland. (1997 UCSC Ph.D. thesis).

Kennedy, Christopher. 2001. Polar opposition and the ontology of 'degrees'. *Linguistics and Philosophy* 24:33–70.

Kennedy, Christopher, and Louise McNally. 1999a. From event structure to scale structure: Degree modification in deverbal adjectives. In *Proceedings from SALT IX*, ed. Tanya Matthews and Devon Strolovitch, 163–180. Ithaca, NY: CLC Publications.

Kennedy, Christopher, and Louise McNally. 1999b. From event structure to scale structure: Degree modification in deverbal adjectives. In *Proceedings from SALT IX*, ed. Tanya Matthews and Devon Strolovitch, 163–180. Ithaca, NY: CLC Publications.

Kennedy, Christopher, and Louise McNally. 2004. Scale structure, degree modification and the semantics of gradable predicates. Unpublished ms.

Klein, Ewan. 1980. A semantics for positive and comparative adjectives. *Linguistics and Philosophy* 4:1–45.

Klein, Ewan. 1991. Comparatives. In *Semantik: Ein internationales handbuch der zeitgenössischen forschung*, ed. Arnim von Stechow and Dieter Wunderlich. Berlin: Walter de Gruyter.

Kyburg, Alice, and Michael Morreau. 2000. Fitting words: Vague language in context. *Linguistics and Philosophy* 23:577–597.

Lasersohn, Peter. 1999. Pragmatic halos. *Language* 75:522–551.

Ludlow, Peter. 1989. Implicit comparison classes. *Linguistics and Philosophy* 12:519–533.

McConnell-Ginet, Sally. 1973. Comparative constructions in English: A syntactic and semantic analysis. Doctoral Dissertation, University of Rochester.

Montague, Richard. 1970. The proper treatment of quantification in ordinary English. In *Richard montague: Selected papers. 1974*, ed. R. Thomason, 247–270. New Haven, Connecticut: Yale University Press.

Paradis, Carita. 2001. Adjectives and boundedness. *Cognitive Linguistics* 12:47–65.

Rotstein, Carmen, and Yoad Winter. to appear. Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics* .

Sapir, Edward. 1944. Grading: A study in semantics. *Philosophy of Science* 11:93–116.

Sedivy, Julie, Michael Tanenhaus, C. Chambers, and Gregory Carlson. 1999. Achieving incremental semantic interpretation through contextual representations. *Cognition* 71:109–147.

Seuren, Pieter A.M. 1973. The comparative. In *Generative grammar in europe*, ed. F. Kiefer and N. Ruwet. Dordrecht: Riedel.

Stanley, Jason. 2000. Context and logical form. *Linguistics and Philosophy* 23.4:391–434.

Stanley, Jason. 2002. Nominal restriction. In *Logical form and language*, ed. G. Peters and G. Preyer, 365–368. Oxford: Oxford University Press.

von Stechow, Arnim. 1984. Comparing semantic theories of comparison. *Journal of Semantics* 3:1–77.

Sybesma, Rynt. 1999. *The mandarin VP*. Boston: Kluwer.

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

Williamson, Timothy. 1994. *Vagueness*. London: Routledge.

Yoon, Yoongeun. 1996. Total and partial predicates and the weak and strong interpretations. *Natural Language Semantics* 4:217–236.

Contact information

Department of Linguistics
Northwestern University
Evanston, IL 60208

kennedy@northwestern.edu