Chris Kennedy
University of Chicago

Seminar on Vagueness
11 April, 2006

The Semantics of Gradable Predicates

1 Introduction

The plan for today is to build a basic understanding of the 'state of the art' in semantic analyses of gradable predicates, so we can begin to understand how sentences like (1) are assigned truth conditions in a context of utterance.

(1) The coffee in Rome is expensive.

Sentences like (1) have three distinguishing characteristics, which have been the focus of a lot of work on vagueness in semantics and the philosophy of language.

TRUTH CONDITIONAL VARIABILITY: (1) could be judged true if asserted as part of a conversation about the cost of living in various Italian cities, as in (2a), but false in a discussion of the cost of living in Chicago vs. Rome, as in (2b).

(2) a. In Rome, even the coffee is expensive!
b. The rents are high in Rome, but at least the coffee is not expensive!

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.

(3) a. Mud Blend: 1.50/kilo
b. Organic Kona: 22.75/kilo
c. House Blend: 11.99/kilo

THE SORITES PARADOX: The structure of the argument in (4) appears to be valid, and the premises appear to be true, but the conclusion is without a doubt false. The problem has to do with the second premise; what is hard is figuring out exactly what goes wrong, and why we don't realize that something is going wrong.

(4) *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.

We'll be spending most of the quarter worrying about borderline cases and the Sorites Paradox; today I want to focus more on the first phenomenon.

2 Degree-based analyses

2.1 The semantics of the positive form

One common set of assumptions about gradable adjective meaning (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 2005a, Kennedy to appear):

(5) 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.

There are various compositional implementations of the core hypotheses about gradable adjective meaning stated in (5); here I will follow Bartsch and Vennemann 1972, 1973 and Kennedy 1999 and analyze gradable adjectives as measure functions (type $\langle e, d\rangle$).

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

This assumption about semantic type not crucial — we could instead assume a $\langle d, \langle e, t\rangle\rangle$ analysis. What will turn out to be crucial are the more general assumptions in (5).

Gradable adjectives are mapped to properties of individuals by DEGREE MORPHOLOGY (comparatives, intensifiers, measure phrases, etc.). For example:

(7) a. $[\![3\ \textit{years}]\!] = \lambda g \in D_{\langle e,d\rangle} \lambda x.g(x) \succeq \textbf{3 years}$
b. $[\![3\ \textit{years}]\!]([\![\textit{old}]\!]) = \lambda x.\textbf{old}(x) \succeq \textbf{3 years}$

To handle the positive (unmarked) form, we have two options:

1. We can assume a null degree morpheme *pos* that introduces the context-dependent meaning associated with this form (see e.g., Bartsch and Vennemann 1972; Cresswell 1977; von Stechow 1984; Kennedy 1999).
2. We can assume a lexical type-shifting rule that turns a measure function into a property of individuals (cf. Neeleman, Van de Koot, and Doetjes 2004; Kennedy and McNally 2005b).

Potential evidence in favor of a morphological analysis comes from Mandarin Chinese (Sybesma 1999):

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

(9) 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'

However, it could be the case that some languages lexicalize what in other languages is a lexical rule (cf. Chierchia 1998). To keep the representations as transparent as possible, I'll assume a null morpheme today, but I think this question is an open one.

- THE LARGER QUESTION: what is the meaning of the positive form morpheme/type-shifting rule?

If gradable adjectives have fixed denotations as measure functions, then the vagueness of the positive form must stem ultimately from properties of *pos*. (Combinations of A plus other degree morphemes are not necessarily vague.)

What we want is something like (10), where $\mathbf{d}_s$ is shorthand for 'a contextually appropriate standard of comparison, whatever that is'.

(10) $[\![pos]\!] = \lambda g \lambda x. g(x) \succeq \mathbf{d}_s$

Most discussions of the positive form stop here, but if we really want to figure out the semantics of these things, we need to push a bit further:

- 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 fixed by extra-linguistic factors?

The standard of comparison can be manipulated in what appears to be a compositional way by constituents local to the predicate:

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

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

(11c) and (12c) can both be true in a situation in which the 'bare' positive in (13a) is false, indicating that they are affecting truth conditions.

(13) 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 (Klein 1980a).

Bartsch and Vennemann (1972), for example, provide a denotation for *pos* that is equivalent to (14), where c is a property and **norm** returns the average degree to which the objects in the set defined by c possess the property measured by g.

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

Contextual variability boils down to the task of finding an appropriate property for the comparison class variable. This is not particularly unusual (and therefore nice).

Additional relevant data: the standard can vary as a function of the value of the argument of the predicate (Kennedy 1999).

(15) 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.

Assume that the comparison class variable can range over properties that are relativized to particular individuals (skolem functions):

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

2.2 Eliminating norms

Unfortunately, reducing the context dependence of the positive form to the identification of a comparison class doesn't help us with borderline cases and the Sorites Paradox (Rusiecki 1985; Pinkal 1995).

- Once the comparison class variable is fixed, the standard of comparison is too: it is the degree on the scale that represents the average degree to which the objects in the comparison class possess the relevant property.

This is too precise, as illustrated by (17).

(17) CONTEXT: we know that the median rent for apartments on the street is $700.
A rent of $750 is expensive for an apartment on this street.

Even though the comparison class is explicit, we might still be unwilling to judge this sentence as true if there are a few apartments with rents significantly higher than $750.

Likewise, the argument in (18) remains paradoxical, even though the generalization in (18P2) should fail for the move from $701 to $700 (Graff 2000).

(18) 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.

- Variability in the comparison class cannot be responsible for vagueness.

Maybe borderline cases and the Sorites Paradox don't have a semantic explanation, but rather an epistemological one (Williamson 1994)? Perhaps, but this doesn't explain why we are willing to accept the second premise as true (Graff 2000).

In fact, the analysis in (14) does not even derive the right truth conditions, and so fails on purely linguistic grounds. As pointed out by Bogusławski (1975), (19) should be a contradiction:

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

This suggests that **norm** cannot have the precise interpretation assumed by Bartsch and Vennemann (1972) (a point they seem to concede). Perhaps a different sort of standard-identifying function will do a better job at deriving vagueness?

This is what I will argue below, but let's first take a closer look at comparison classes. It turns out that we can get rid of them too.

2.3 Eliminating comparison classes

The facts I presented above seemed to indicate that modified nominals and *for*-PPs can compositionally affect the computation of the standard of comparison.

1. Although modified nominals may provide the basis for computing a standard of comparison, they do not have to:

(20) 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.

(21) 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.

There is nothing compositional about the relation between a modified nominal and the standard, and nothing obligatory (contra Wheeler 1972), so:

- Modified nominals provide no evidence for a comparison class argument of *pos*.

2. A predicate of the form *A for a NP* modifier introduces the presupposition that the argument of A is an NP (cf. Wheeler 1972, p. 316).

(22) 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?

(23) 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?

(24a) 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 (24b).

(24) 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.

Modified nominals do not introduce this presupposition, even when they provide the comparison class.

(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?

(26) a. ??That elephant is (really) large for a mouse.
b. That elephant is (really) a large mouse.

The semantics of the positive form in (14) assigns a denotation along the lines of (27) to the adjectival predicates in (26). Clearly, there is no presupposition that the argument of the predicate be a mouse, which is good for (26b) but bad for (26a).

(27) $\lambda x.\mathbf{large}(x) \succ \mathbf{the\ size\ of\ the\ average\ mouse}$

We could augment the analysis by stipluating the presupposition, but this would then be good for (26a) but bad for (26b).

MY CONCLUSION: A *for*-PP 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 hypothesis is implemented in (28) (I assume that the NP contributes a property, though it could be a set or a kind; see Graff 2000 for relevant discussion).

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

If *expensive* is a function from objects to degrees of cost, then *expensive for a Honda* is a function from Hondas to degrees of cost:

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

The examples in (23) are anomalous because the semantic argument of the adjective (the subject) is not a member of its domain. The modified nominals in (25) do not restrict the domains of the adjectives that modify them, so similar effects do not arise.

But if *for*-PPs are domain restrictions on gradable adjectives, they cannot also serve as comparison class arguments to *pos*. Therefore:

- *For*-PPs provide no evidence for a comparison class argument of *pos*.

3. If a *for*-PP can explicitly restrict the domain of a gradable adjective, then it is reasonable to assume that the domain can be implicitly restricted as well. Let us assume that there is a null *for*-PP that does this job (cf. von Fintel 1994; Stanley 2000; Stanley and Szabó 2000; Martí 2002; Giannakidou 2004 on quantifier domain restrictions):

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

Here **f** is a variable over properties; skolemization gets us the binding facts:

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

The truth conditions of (31) are: every member of my family x is tall wrt to a domain determined by x.

We can also handle 'sloppy' domains in ellipsis cases quite straightforwardly (Klein 1980a; Ludlow 1989):

(32) a. That elephant is large, and that flea is too.
b. $\underline{[\![pos]\!](\lambda y : \mathbf{f}(y).\textbf{large}(y))}(\textit{that elephant}) \wedge \underline{[\![pos]\!](\lambda z : \mathbf{f}(z).\textbf{large}(z))}(\textit{that flea})$

Since we don't need a comparison class to derive these results, we can conclude that:

- 'Bound comparison class' interpretations of the positive form provide no evidence for a comparison class argument of *pos*.

Conclusion: Comparison classes are not semantic (or syntactic) arguments of the positive form. Instead, they are simply discourse salient properties used to compute the standard of comparison. *Used by what?*

2.4 A vague semantics for the positive form

Following Graff 2000 and Bogusławski 1975, I propose the semantic analysis of *pos* stated in (33), where **s** is a context dependent function from adjective denotations to degrees that returns the minimum degree on the adjective's scale that represents a **significant** amount of the measured property in the context of utterance.

(33) $[\![pos]\!] = \lambda g \lambda x.g(x) \succeq \mathbf{s}(g)$

The **s** function in (33) corresponds to the DELINEATION FUNCTION of Lewis 1970 and Barker 2002, which is a contextual coordinate that maps a gradable adjective meaning to a degree. The difference is the addition of the 'significance' requirement, which has some important consequences. In particular, it gets us:

Truth conditional variability Significance is an inherently relative notion: in order to figure out what counts as a significant degree of e.g. height, we need to restrict our attention to a particular domain (significant relative to 3 year olds, etc.). Because **s** is a context sensitive function, it may look to the discourse model to find a salient property on which to calculate significance. These properties are what we call 'comparison classes', but they do not have representational status in the sentence/proposition.

Borderline cases As pointed out by Graff, even if we know the domain on which to base a judgment of significance, we may have incomplete knowledge of other relevant factors (the task at hand, the interests/expectations of the participants, etc.). We may therefore not be able to pinpoint the actual standard, resulting in borderline cases.

The Sorites Paradox Let us assume with Graff a 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 (cf. Kamp 1981; Soames 1999).

We ought to be able to derive this from the status of **s** as a context-sensitive function, with an appropriate dynamic framework for the interpretation of vague expressions (see e.g. Kyburg and Morreau 2000; Barker 2002): 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).

According to Graff, this explains why we are inclined to accept the second premise of a Sorites argument as true, even if it is strictly speaking false. The universal generalization would require evaluation of pairs that fall under the Similarity Constraint. But the act of examining two such objects to see if the standard of comparison is between them has the consequence that the standard cannot be between them!

We also derive the fact that (19) is not contradictory: it is possible to have a greater than average degree of some property without having a significant degree.

But is the 'significant degree' requirement really part of the conventional meaning (the content of the **s** function) of the positive form, or does it come from some other, non-semantic property of the use of vague/imprecise terms?

Evidence that it is a matter of semantics comes from the fact that its effects can be seen even in examples that involve maximally explicit standards of comparison, such as (34).

(34) This novel is long compared to that one (though both are quite short).

(34) 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*.

1. Assume that the *compared to* phrase causes the predicate to be evaluated with respect to a context that includes only the two objects being compared.

2. Assume that the standard of comparison must partition the domain of the predicate into two non-empty sets (a 'positive' and 'negative' extension; Klein 1980a).

If the positive form is true of the novel it is being predicated of, it must be false of the second novel. This etnails that the first is longer than the second.

What is relevant to the current discussion is the fact that this use of the positive form is infelicitious in contexts requiring CRISP JUDGMENTS:

(35) 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.

This is expected given the assumptions above and the semantics of *pos* in (33):

- The context in (35a) is one in which only the two novels 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.

If the truth conditions of the positive form did not include the 'significant degree' requirement, this explanation would be unavailable.

In a context such as (35), we could satisfy the domain partitioning constraint by setting the standard to the degree of length of the shorter novel, deriving truth conditions equivalent to those of the comparative form.

2.5 Domain dependence

The hypothesis that the standard function **s** is a function from gradable adjectives to degrees provides a basis for explaining the semantic function of *for*-PPs.

If a *for*-PP restricts the domain of the adjective, then the **s** function has different inputs in (36a) and (37a).

(36) a. This car is expensive for a Honda.
 b. $\lambda x : \mathbf{Honda}(x).\mathbf{expensive}(x) \succ \mathbf{s}(\lambda y : \mathbf{Honda}(y).\mathbf{expensive}(y))$

(37) a. This car is expensive for a Fiat.
 b. $\lambda x : \mathbf{Fiat}(x).\mathbf{expensive}(x) \succ \mathbf{s}(\lambda y : \mathbf{Fiat}(y).\mathbf{expensive}(y))$

Since the domains of the adjectives are different, the functions they express are different, and as a result, the values returned by **s** could be different.

However, it is not clear that this analysis captures the fact that a restricted adjectival domain must be used to compute the standard of comparison, a fact illustrated by the following example.

CONTEXT: an experiment is being run involving a subject who watches insects moving around a maze and describes what is happening. Sometimes insects leave the maze, sometimes more insects enter; sometimes other creatures enter the maze as well. Suddenly, a rat-like creature of substantial size (especially compared to the insects) enters the maze. The subject first utters (38), then one of the continuations in (38a-c).

(38) Another animal just entered the maze.
 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.

(38a-b) are both fine, but (38c) is infelicitous. In particular, it is obviously false: even a rat (or a rat-like creature) of subtantial size is not large for a mammal.

- A restricted adjectival domain must be used to calculate the standard of comparison, even in a context in which other properties are both highly salient and derive more felicitous truth conditions.

Does this follow? I don't think so. In particular, the inherent context dependence of **s** (the need to find a propety in the context on which to calculate significance) seems to allow for the possibility that some salient property other than the domain restriction on the adjective could be used to calculate the standard, especially if using the domain restriction alone resulted in an infelicitous interpretation, as in (38c).

So we seem to have evicence for a (descriptive) constraint along the lines of (39).

(39) *Domain Dependence*
If a gradable adjective has a restricted domain, then the standard of comparison must be calculated relative to the domain.

A FUNCTIONAL EXPLANATION: Why bother to restrict the domain if not to specify the basis on which the standard of comparison should be computed?

Not implausible, but I think there's a deeper — and more grammatical — principle at work here. But to see what this principle is, we need to broaden our empirical coverage and turn to an examination of absolute gradable adjectives.

3 Getting by without degrees

A very different sort of approach to the semantic analysis of gradable adjectives is developed by Klein (1980b) (see also McConnell-Ginet 1973). This is essentially a linguistic implementation of a supervaluation analysis (Kamp 1975; Fine 1975), which we will be discussing in a lot more detail next week, so I am not going to go too deeply into the technical details now. Instead I will focus on the general form of the analysis, and we will consider its broader implications for the analysis of vagueness next week.

Klein claims that a semantic theory of GAs must satisfy (40) to be explanatorily adequate.

(40) *Principle of Compositionality*
If A is a gradable adjective, then the meaning of [Aer than __] is a function of the meaning of A.

Klein claims that degree analyses fail to satisfy this criterion, because they define the meaning of GAs in terms of a comparison relation ($\succeq$). Is this a valid criticism?

- Sapir (1944): Comparison is a psychological primitive. If this is correct, and if we can link a degree-based anlysis to a general theory of comparison/measurement, then Klein's objection is at least partially dealt with.
- If GAs denote measure functions, then the comparative form — and the positive form too — *is* a function of the meaning of A, which does not itself encode comparison.

At any rate, Klein's strategy for satisfying (40) is to begin from the assumption that GAs have the same semantic type as other predicates — they are functions from objects to truth values. They differ in two crucial ways, however:

1. They are partial functions from U to $\{0,1\}$. That is, a gradable adjective G partitions its domain into three sets:

(a) the positive extension of G: $\{x \mid [\![G(x)]\!] = 1\}$

(b) the negative extension of G: $\{x \mid [\![G(x)]\!] = 0\}$

(c) the extension gap of G: $\{x \mid [\![G(x)]\!]$ is undefined$\}$

2. Their interpretations are context-dependent, such that the exact compositions of the positive and negative extensions and extension gap are a function of the context of utterance.

Specifically, Klein assumes that every context c determines a *comparison class* of objects that provides the domain of the adjective. This allows him to maintain the assumption that e.g., *expensive* has a single core meaning along the lines of (41):

(41) For any context c, comparison class $\kappa(c)$, and object a:

a. $[\![expensive(a)]\!]_c^\kappa = 1$ iff a is definitely expensive,

b. $[\![expensive(a)]\!]_c^\kappa = 0$ if a is definitely not expensive, and

c. $[\![expensive(a)]\!]_c^\kappa$ is undefined otherwise.

The variability comes from the fact that the composition of the comparison class will affect what does and does not 'count as' definitely expensive.

(42) The Mars Pathfinder mission was expensive.

a. $\kappa(c_1) = \{x \mid x$ is named 'Pathfinder'$\}$

b. $[\![expensive(p)]\!]_{c1}^\kappa = 1$

c. $\kappa(c_2) = \{x \mid x$ is a mission to outer space$\}$

d. $[\![expensive(p)]\!]_{c2}^\kappa = 0$

But what decides whether something is 'definitely expensive' or not? Note that there is nothing in Klein's representations that corresponds to the standard of comparison — this is evidently outside the linguistic system. (This is an important difference from the degree analysis, where the answer to this question is 'greater than the standard of comparison'.) When we look into supervaluation analyses in more detail next week, we will see what kind of answer is given to this question.

Digression 1: Why not represent the comparison class as a free variable in the lf?

(43) $[\![expensive(\kappa)(pathfinder)]\!]_c$

Klein argues that ellipsis constructions provide evidence against treating the comparison class (and by extension, the standard of comparison in a degree analysis) as a free variable.

In general, deictically identified expressions maintain their reference across ellipsis:

(44) a. Jorge didn't ride it, but Ivan did. (same 'it')

b. Jorge rode a big one, and Ivan did too. (same kind of big thing)

This need not be true of comparison classes:

(45) Jorge's elephant was big, and Ivan's burro was too.

Klein doesn't say why contextual parameters are <u>not</u> fixed under ellipsis (presumably because they're not part of the linguistic representation), but he makes the interesting observation that the behavior of comparison classes (or standards of comparison) in ellipsis is similar to what we see with quantifier domain restrictions:

(46) Leo gave a bridge party at home yesterday and Jude took the kids swimming. Leo thought everyone had a good time, and Jude did too.

In a degree analysis, these facts point to the quantificational analysis of the standard of comparison that we saw above.

Digression 2: According to Klein, the comparison class provides the domain for the adjective. It follows that the subject should be presupposed to be a member of the comparison class. As we have seen, this is correct:

(47) a. ??CK is short for an NBA player.
b. CK would be short for an NBA player.

However, this is stipulated on Klein's analysis: nothing about the architecture of the theory forces this to be the case. We could fix this by doing what we did above: change the comparison class function to a 'domain restrictor' function and assume that the extensions of the predicate are fixed relative to its domain.

Comparatives: Klein's analysis starts from his analysis of degree modifiers in which the function of a degree modifier is to provide a new comparison class for the adjective it modifies. For example, *very expensive* has the meaning in (48).

(48) $[\![very(expensive)]\!]_c^{kappa} = [\![expensive]\!]_{c[X]}^{\kappa}$, where X is the positive extension of *expensive* at c.

In other words, *very expensive* is a vague predicate just like *expensive*, except that the comparison class for the former is just the positive extension of the latter.

- The analysis of degree modifiers suggests a straightforward explanation of examples like *Julian is a short tall boy*, *Julian is a tall tall boy*, etc., or stacked modifiers *Julian is very very tall.*

Comparatives involve existential quantification over degree modifier meanings:

(49) a. The Viking mission was more expensive than the Pathfinder mission.
b. $\exists d[[\![d(expensive)(v)]\!]_c^{\kappa} = 1 \wedge [\![d(expensive)(p)]\!]_c^{\kappa} = 0]$

In other words, (49a) is true in a context c just in case there is some way of fixing the comparison class for *expensive* such that *The Viking mission was expensive* is true wrt that c-class and *The Pathfinder mission was expensive* is false wrt that c-class.

Since the derived c-class need not be the same as the one determined by the context of utterance, we will not get entailments from the comparative form to the unmarked form. We have some questions, though:

1. This analysis of *very* seems to predict that *very expensive* presupposes that its argument is expensive. Is this correct?
2. What does the analysis predict about things like ??*Julian is more very tall than Nigel*?
3. Since this analysis crucially does <u>not</u> introduce scales and degrees into the linguistic system — for Klein, *there is no 'semantics of degree'!* — we should not expect to find phenomena that are crucially sensitive to properties of these objects. We should however expect to find phenomena that are crucially sensitive to (properties of) the comparison class (though we probably also expect this on a degree analysis).

4. More generally, since the analysis of comparatives is built on top of the basic mechanism for handling vagueness (context-dependent specification of the positive/negative extensions of a predicate relative to its domain/comparison class), we expect all (and only) gradable predicates to be vague.

In fact, the last point is true of the degree analysis as well. A problem for both accounts is that this prediction fails, though as we will see, only a degree-based analysis provides the basis for an explanation of when and why.

4 Absolute gradable adjectives

4.1 The problem

A prediction of both analyses of the positive form that we have considered is that any gradable adjective — any expression of type $\langle e, d \rangle$, or any partial function with context-sensitive, domain dependent extension — should should show the properties of vagueness in the positive form, because these properties stem from the semantics of the positive degree morpheme *pos* (in the degree analysis) and the basic meaning of such predicates (in a Klein-style analysis).

This prediction is incorrect: ABSOLUTE gradable adjectives are demonstrably gradable but do not have context dependent interpretations, do not give rise to borderline cases, and do not trigger the Sorites Paradox (Unger 1975; Rotstein and Winter 2004; Kennedy and McNally 2005a).

MINIMUM STANDARD absolute adjectives merely require their arguments to possess some minimal degree of the property they describe:

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

MAXIMUM STANDARD absolute adjectives such as those in (51) require their arguments to posses a maximal degree of the property in question.

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

Evidence that absolute adjectives are gradable comes from the fact that they are perfectly acceptable in comparatives and with other degree morphology:

(52) 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.

(53) 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:

(54) 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 PROBLEM: if absolute adjectives have the same semantic type as relative adjectives, their interpretations in the positive form are unexpected. Here I'll just focus on the degree analysis; it's easy to reproduce the problem in a Klein-style analysis.

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

(55a) is true of an object if it has a contextually significant degree of impurity, and (55b) is true of an object if it has a contextually significant degree of purity. These are not accurate characterizations of the truth conditions of these predicates.

Furthermore, absolute adjectives should be just as vague as their relative counterparts. As the next sections will demonstrate in detail, this prediction is incorrect.

4.2 Empirical support for the absolute/relative distinction

4.2.1 Imprecision vs. vagueness

Maybe the problem is with my empirical claims, not with the theory. (56) can be felicitously used to describe a situation in which only a very few people show up to a film in a very large movie theater.

(56) The theater is empty tonight.

MY RESPONSE: This is IMPRECISION, not vagueness (Pinkal 1995). Imprecision is most clearly illustrated by measure phrases:

(57) a. The rod is 10 meters long.
b. The rod is long.

Both (57a) and (57b) could be 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.

However, (57a) but not (57b) allows for what Pinkal (1995, pp. 99-100) calls NATURAL PRECISIFICATIONS: we can construct a context in which *10 meters long* distinguishes between objects based on 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 *long*.

(58) 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.

Absolute adjectives allow natural precisifications. If I am a detective in search of a violent criminal, and I think he is in the movie theater described above, I would judge a projectionist who responds to my question *Is anyone in the theater tonight?* with (56) to be lying.

Likewise, (59) is perfectly natural in a spacecraft construction context, even though it implies that a tiny bend is enough to prevent the rod from counting as straight.

(59) 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.

If the basic meanings of *empty*, *straight*, etc. are as I claimed above, then they should indeed allow for the same sort of precisifications as measure phrases.

4.2.2 Entailments

If what I said earlier is correct, then the truth conditions of the positive forms of minimum and maximum standard absolute adjectives are as in (60a-b), respectively.

(60) a. $\lambda x.g_{min}(x) \succ min(\textsc{scale}(g_{min}))$
b. $\lambda x.g_{max}(x) = max(\textsc{scale}(g_{max}))$

(60a) predicts that a negative assertion x *is not* A_{min} should entail that x possesses no amount *adj*-ness at all (assuming that the minimal degree on a scale represents a zero amount of the relevant property). This is correct:

(61) 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.

(60b) predicts that an assertion of x *is* A_{max} should entail that x cannot be more A than it is. This is difficult to test, since maximum standard adjectives allow imprecision. But we can force a precise interpretation by adding focal stress (specifically, a falling tone; cf. Unger 1975), and indeed the expected entailments arise:

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

Neither of the above entailments hold for relative adjectives. Negation is compatible with a positivedegree of the measured property, and assertion does not rule out higher values:

(63) a. Sam is not tall, but his height is normal for his age.
b. That film is interesting, but it could be more interesting.

A related argument involves pairs of antonyms such that negation of one form entails the assertion of the other (Cruse 1986; Rotstein and Winter 2004):

(64) 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.

Both members of the pairs in (64) are absolute adjectives, but the positives have minimum standards while the negatives have maximum standards. Since a minimal positive degree corresponds to a maximal negative degree on the same scale (see Kennedy 2001), the entailments follow from the truth conditions in (60).

Relative antonyms do not show the same entailment relations, because there is no necessary corrleation between their standards — they are not linked by scale structure.

(65) 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.

Finally, comparatives with absolute adjectives generate positive and negative entailments to the positive form, respectively, depending on whether we have a minimum or maximum standard adjective (cf. Rusiecki 1985):

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

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

(66a) is true only if the floor has some degree of wetness: if it had zero wetness, then it could not be wetter than the countertop. This satisfies (60a), entailing (66b).

Similarly, in order for (67a) to be true, it must be the case that the countertop is not maximally dry. It therefore follows from (60b) that the countertop is not dry.

A canonical property of comparatives with relative adjectives, is that they do not give rise to positive or negative entailments in the comparative form:

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

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

This is expected: the 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.

4.2.3 The Sorites Paradox

As we have seen, a defining characteristic of relative adjectives is that they give rise to the Sorites Paradox. This is illustrated again for the relative adjective *big* in (70).

(70) 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.

Following Graff 2000, I have argued that the paradox arises from the 'contextually significant' semantics of the positive form of a relative adjective. If absolute adjectives had the same meanings in the positive form, they should also give rise to the Paradox, but this is not the case.

In (71), for example, the second premise is quite naturally judged to be false, bypassing the paradox.

(71) 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.

(72) shows the same thing for the minimum standard adjective *impure*.

(72) 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.

P2 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!

4.2.4 Summary

The standard of comparison of the positive form of an absolute adjective a minimum or maximum degree (at least as a default); it is not a contextually significant degree of the gradable property expressed by the adjective.

This is unexpected on both analyses of the positive form under consideration. In the rest of this handout, I will show how we can account for the facts within a degree-based analysis; we will take up the question of whether a Klein-style analysis can accommodate the facts in detail next week.

5 Vagueness and grammar

5.1 The polysemy of the positive form

The descriptive conclusion to draw from the facts discussed in the preceding sections is that the positive form is polysemous: it can have (at least) the three distinct interpretations specified in (73), depending on the kind of adjective that it combines with.

(73) $[\![pos]\!] =$

a. $\lambda g \lambda x. g(x) \succeq \mathbf{s}(g)$ (when g is a relative A)
b. $\lambda g \lambda x. g(x) \succ min(\text{SCALE}(g))$ (when g is an absolute minimum A)
c. $\lambda g \lambda x. g(x) = max(\text{SCALE}(g))$ (when g is an absolute maxumum A)

This raises the following questions:

1. What is the 'core meaning' of the positive form?
2. How can we explain (in a principled way) the correlation between particular adjectives and particular interpretations of *pos*?

I think the answer to the first question is to some extent a functional one: the meaning of the positive form is a mapping from a gradable adjective (a measure function) to a property of individuals that relates 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).

There are several natural options for establishing such a reference point:

- The maximal and minimal degrees on a scale are fixed values that 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 supports a meaningful partitioning of the domain of the adjective.
- Others...??? (Color terms, *early/late*, 'extreme' adjectives, etc.)

We still have to deal with the second question. For example, why does *long* combine only with (73a), deriving a relative interpretation, while *bent* combines only with (73b) and *straight* only with (73c), deriving minimum and maximum absolute interpretations, respectively? Why don't absolute adjectives like *straight* and *bent* also combine with (73a), deriving relative interpretations? More generally:

- What factors are responsible for the fact that relative adjectives give rise to vague interpretations in the positive form, but absolute adjectives do not?

5.2 Scale structure and standard of comparison

If gradable adjectives denote functions that map objects onto scales (or contain such functions as part of their meanings), they could differ with respect to properties of the scalar representations they make use of.

Kennedy and McNally (2005a) argue that one parameter of scalar variation is the open/closed distinction: whether a scale has minimal or maximal elements (see also Kennedy and McNally 1999; Paradis 2001; Rotstein and Winter 2004).

(74) *A typology of scale strucures*

a. (TOTALLY) OPEN: ○————————○

b. LOWER CLOSED: ●————————○

c. UPPER CLOSED: ○————————●

d. (TOTALLY) CLOSED: ●————————●

Evidence for this typology comes from the distribution of degree modifiers that pick out maximal/ minimal degrees (Rotstein and Winter 2004; Kennedy and McNally 2005a).

The reasoning relies the observation that antonymous pairs of gradable adjectives map their arguments onto the same scale, but impose inverse orderings on their shared domains (Sapir 1944; Seuren 1978; Rullmann 1995; Kennedy 2001).

- 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).
- Negative adjectives measure decreasing amounts of a property (if a is shorter than b, then a has less height than b).

This leads to the following predictions:

1. 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 and the range of the negative member should include a maximum degree.
2. If a scale is closed on the upper end, then the postive antonym should include a maximum degree and the range of the negative antonym should include a minimum degree.
3. 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.
4. If a scale is open on the upper end, the positive should have no maximum and the negative no minimum.

The predictions for degree modifiers that pick out maximal (e.g., *absolutely, completely, totally, perfectly*) or minimal (e.g., *slightly, partially*) degrees:

(75)

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

Perfectly and *slightly* provide particularly clear judgments across a range of cases:

(76) *Open scales*
a. ??perfectly/??slightly {tall, deep, expensive, likely}
b. ??perfectly/??slightly {short, shallow, inexpensive, unlikely}

(77) *Lower closed scales*
a. ??perfectly/slightly {bent, bumpy, dirty, worried}
b. perfectly/??slightly {straight, flat, clean, unworried}

(78) *Upper closed scales*
a. perfectly/??slightly {certain, safe, pure, accurate}
b. ??perfectly/slightly {uncertain, dangerous, impure, inaccurate}

(79) *Closed scales*
a. perfectly/slightly {full, open, opaque}
b. perfectly/slightly {empty, closed, transparent}

The crucial generalizations (Kennedy and McNally 2005a):

(80) a. Gradable adjectives that use open scales have relative interpretations in the postive form.
b. Gradable adjectives that use closed scales have absolute interpretations in the positive form.

(80a) follows directly from the interaction of scale structure and the polysemous semantics of *pos*: open scales lack maximal and minimal values, so open scale adjectives are incompatible with maximum and minimum standards, leaving only the relative interpretation.

(80b), however, does not follow. A closed scale structure may restrict the range of interpretations that *pos* can have, but it does not uniquely determine a particular interpretation. In particular, there is no semantic incompability between (totally or partially) closed scales and a relative interpretation of *pos*.

The possibility that such interpretations can be forced in specific contexts illustrates this, but the robust and systematic evidence that closed scale adjectives have default absolute interpretations indicates that some other constraint, sensitive to this particular feature of gradable adjective meaning, is constraining the interpretation of the positive form.

5.3 Interpretive economy

We have already seen evidence that features of the conventional meaning of a gradable adjective can have consequences for the truth conditions of the positive form:

(81) *Domain Dependence*
If a gradable adjective has a restricted domain, then the standard of comparison must be computed 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'):

(82) 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:

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

The basic idea is that truth conditions should 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 as a last resort, when conventional meaning is insufficient.

I need to say something about what kind of constraint this is, but first I will provide empirical support for it by showing how it helps explain the data I have presented today.

5.3.1 Restricted Domains

Consider first the case of open scale adjectives with unrestricted vs. restricted domains.

(84) 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 interpretations of the predicates:

(85) 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 inputs to **s** are different functions, so the standards are allowed to differ. *Interpretive Economy* ensures that they will differ:

- Significance is an inherently relative notion: in order to determine the minimal degree that represents a significant amount of size, you need a property on which to base this judgment.
- The conventional meaning of the adjective in (85a) is a function from objects that have sizes to sizes. The mere property of 'having size' 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 a likely candidate, though as we have seen this is only a preference, not a requirement.
- The conventional meaning of the (modified) adjective in (85b) is a function from mammals 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.

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). But use of the restricted domain in (85b) maximizes the contribution of the conventional meaning of the constituents of the predicate to the calculation of its truth conditions, in accord with *Interpretive Economy*.

5.3.2 Partially Closed scales

For adjectives with closed scales, *Interpretive Economy* chooses between the competing senses of *pos*, favoring an absolute interpretation.

Consider the adjectives *bent* and *straight*, which have lower-closed and upper-closed scales, respectively. For *bent*, both of the interpretations in (86) are in principle possible.

(86) a. $\lambda x.\mathbf{bent}(x) \succeq \mathbf{s}(\lambda y.\mathbf{bent}(y))$
b. $\lambda x.\mathbf{bent}(x) \succ min(\text{SCALE}(\lambda y.\mathbf{bent}(y)))$

The truth conditions of (86b) are computed strictly on the basis of the conventional meaning of *bent* (its range), while (86a) introduces context dependence via the **s** function. *Interpretive Economy* therefore selects the absolute interpretation in (86b).

For *straight*, the choice is between the relative interpretation in (87a) and the maximum standard intepretation in (87b).

(87) a. $\lambda x.\mathbf{straight}(x) \succeq \mathbf{s}(\lambda y.\mathbf{straight}(y))$
b. $\lambda x.\mathbf{straight}(x) = max(\text{SCALE}(\lambda y.\mathbf{straight}(y)))$

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.

5.3.3 Totally Closed Scales

Interpretive Economy predicts that totally closed adjectives should have absolute interpretations for the same reasons outlined above, but it doesn't actually choose between the minimum and maximum standard interpretations of *pos*. Since both of these are compatible with totally closed scales, we might expect such adjectives with such scale structures to give rise to interpretive variability in the positive form.

The antonyms *opaque* and *transparent* verify this prediction. (88a) 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).

(88) CONTEXT: Manipulating a device that changes the degree of tint of a car window from 0% (completely transparent) to 100% (completely opaque)
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, (88b) 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 (complete transparency) and *opaque* has a minimum standard (partial opacity).

There is definitely a preference for maximum standard interpretations of these adjectives, but this can be explained in pragmatic terms: a maximum standard interpretation entails a minimum standard one, but not vice-versa.

Other total scale adjectives that allow both minumum and maximum standard intepretations are *open, exposed* and *uncovered*.

(89) CONTEXT: On the starship *Enterprise*
  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.

Anyone who does not understand (89a) to be a warning that partial opening of the airlock will result in depressurization, is a danger to the ship and crew, and so is a helmsman who fails to understand (89b) as an prohibition against trying to leave the station before the space door is completely open.

Similarly, (90a) 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).

(90) a. an exposed line of troops
  b. Janet Jackson's exposed breast

However, the reason that ABC got in trouble with the FCC over (90b) (during the halftime show of Superbowl XXXVIII) was because Janet Jackson's breast was exposed all the way to the nipple, not because there was some skin visible.

Despite these examples, however, many (if not most) totally closed scale adjectives have fixed standards of comparison:

(91) a. closed, hidden, covered
  b. full/empty

Is the difference simply a matter of lexical idiosyncrasy? That would to some extent undermine the whole enterprise....

In fact, the choice of standard is systematic. Focusing on the case of deverbal adjectives derived from accomplishment verbs, Kennedy and McNally (2005a) observe that the standard associated with a total scale deverbal adjective 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 the incremental theme of the source verb, then the adjectival form has a maximum standard.
- If not, then the adjectival form has a minimum standard.

For illustration, consider *loaded*. (92a-b) show that a goal argument gets a minimum standard, but an incremental theme argument gets a maximum standard.

(92) 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 show that the positive form of a deverbal adjective entails the completion of an eventuality corresponding to the one described by the source verb. This leads to the following predictions:

- In the case of an incremental theme argument, 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. A minimum standard interpretation is therefore filtered out.
- With a non-incremental theme argument, 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. Either a minimum or maximum standard interpretation should be possible.

In fact, *loaded* can also have a maximum standard interpretation when it has a goal argument:

(93) 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.

Most adjectives with totally closed scales appear to be deverbal, though not all are related to accomplishment verbs. The adjectives in (94), for example, are related to achievements, but the core of the explanation outlined above carries over directly.

(94) a. open, uncovered, exposed
b. closed, covered, hidden

The verbal source of the adjectives in (94a) (*open, uncover, expose*) names the initiation point of an event, while the verbal source of those in (94b) (*close, cover, hide*) names the culmination point.

- An argument of the 'initiation predicate' counts as having participated in the event named by the verb as long as it has a non-zero degree of the measured property; this is consistent with a minimal or maximal standard.
- An argument of the 'culimation predicate' counts as having 'participated in the event just in case it has a maximal degree of the relevant property; this is consistent only with a maximal standard.

For lexical adjectives like *full/empty* vs. *opaque/transparent*, we need to assume that the former but not the latter are related to events in the same way as deverbal adjectives.

Support for this comes from derivational morphology, and from the fact that *full/empty* are stage level (in the positive form), while *opaque/transparent* are individual level:

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

6 Concluding Remarks

A central feature of the analysis developed here is that the vagueness of the positive form of a gradable adjective is not a consquence of gradability per se, but is rather contingent upon a formal property of the representation of gradability in the semantics: the structure of an adjectival scale.

If the analysis is correct, then we can draw the following conclusions:

- At least in the case of gradable predicates, whether an expression is vague or not depends on its grammatical (lexical semantic) properties.
- The semantics of gradable predicates must be stated in terms of scalar representations.

These results are important because they call into question approaches to the semantics of gradable predicates that essentially take the opposite strategy from the one advocated here, by deriving gradability from a general theory of vagueness, eliminating the need for a sematnic ontology that includes scales and degrees. That is the essence of Klein's supervaluation-based analysis. But to see whether the absolute/relative distinction is a problem for such accounts in general (as opposed to just Klein's implementation), we need to take a closer look at how they work. This will be the topic of next week's discussion.

References

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

Bartsch, Renate, and Theo Vennemann. 1972. The grammar of relative adjectives and comparison. *Linguistische Berichte* 20:19–32.

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.

Bogusławski, Andrzej. 1975. Measures are measures: In defence of the diversity of comparatives and positives. *Linguistiche Berichte* 36:1–9.

Chierchia, Gennaro. 1998. Reference to kinds across languages. *Natural Language Semantics* 6:339–405.

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

Cruse, D. Allan. 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.

Giannakidou, Anastasia. 2004. Domain restriction and the arguments of quantificational determiners. In *Proceedings of SALT 14*, ed. Editor, xx–xx. Ithaca, NY: CLC Publications.

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

Heim, Irene. 1985. Notes on comparatives and related matters. Ms., University of Texas, Austin, TX.

Heim, Irene. 2000. Degree operators and scope. In *Semantics and Linguistic Theory 10*, ed. Brendan Jackson and Tanya Matthews, 40–64. Ithaca, NY: CLC Publications.

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

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

Kamp, Hans. 1981. The paradox of the heap. In *Aspects of philosophical logic*, ed. U. Mönnich. Dordrecht: Reidel.

Kennedy, Christopher. 1999. *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. to appear. Vagueness and grammar: The semantics of relative and absolute gradable predicates. *Linguistics and Philosophy*.

Kennedy, Christopher, and Louise McNally. 1999. From event structure to scale structure: Degree modification in deverbal adjectives. In *Semantics and Linguistic Theory 9*, ed. Tanya Matthews and Devon Strolovitch, 163–180. Ithaca, NY: CLC Publications.

Kennedy, Christopher, and Louise McNally. 2005a. Scale structure and the semantic typology of gradable predicates. *Language* 81:345–381.

Kennedy, Christopher, and Louise McNally. 2005b. The syntax and semantics of multiple degree modification in english. In *The Proceedings of the 12th International Conference on Head-Driven Phrase Structure Grammar, Department of Informatics, University of Lisbon*, ed. Stefan Müller, 178–191. Stanford: CSLI Publications.

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

Klein, Ewan. 1980b. 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, 673–691. Berlin: Walter de Gruyter.

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

Lewis, David K. 1970. General semantics. *Synthese* 22:18–67.

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

Martí, Maria Luisa. 2002. Contextual variables. Doctoral Dissertation, University of Connecticut.

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

Neeleman, Ad, Hans Van de Koot, and Jenny Doetjes. 2004. Degree expressions. *The Linguistic Review* 21:1–66.

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

Pinkal, Manfred. 1995. *Logic and lexicon*. Dordrecht: Kluwer.

Rotstein, Carmen, and Yoad Winter. 2004. Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics* 12:259–288.

Rullmann, Hotze. 1995. Maximality in the semantics of wh-constructions. Doctoral Dissertation, University of Massachusetts, Amherst, MA.

Rusiecki, Jan. 1985. *On adjectives and comparison in English*. New York: Longman guistics Library.

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

Seuren, Pieter A.M. 1973. The comparative. In *Generative grammar in Europe*, ed. Ferenc Kiefer and Nicolas Ruwet, 528–564. Dordrecht: Riedel.

Seuren, Pieter A.M. 1978. The structure and selection of positive and negative gradable adjectives. In *Papers from the Parasession on the Lexicon, Chicago Linguistics Society*, ed. W.M. Jacobsen Farkas, Donka and K.W. Todrys, 336–346. Chicago: Chicago Linguistic Society.

Soames, Scott. 1999. *Understanding truth*. Oxford: Oxford University Press.

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

Stanley, Jason, and Zoltan Szabó. 2000. On quantifier domain restriction. *Mind and Language* 15:219–261.

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.

Wheeler, Samuel. 1972. Attributives and their modifiers. *Noûs* 6:310–334.

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