

Vagueness and grammar: the semantics of relative and absolute gradable adjectives*

Christopher Kennedy

Received: 23 August 2006/Accepted: 15 September 2006
© Springer Science+Business Media B.V. 2007

Abstract This paper investigates the way that lexical semantic properties of linguistic expressions influence vagueness, focusing on the interpretation of the positive (unmarked) form of gradable adjectives. I begin by developing a semantic analysis of the positive form of ‘relative’ gradable adjectives, expanding on previous proposals by further motivating a semantic basis for vagueness and by precisely identifying and characterizing the division of labor between the compositional and contextual aspects of its interpretation. I then introduce a challenge to the analysis from the class of ‘absolute’ gradable adjectives: adjectives that are demonstrably gradable, but which have positive forms that relate objects to maximal or minimal degrees, and do not give rise to vagueness. I argue that the truth conditional difference between relative and absolute adjectives in the positive form stems from the interaction of lexical semantic properties of gradable adjectives—the structure of the scales they use—and a general constraint on interpretive economy that requires truth conditions to be computed on

* This paper grew out of joint work with Louise McNally, and owes a great debt to this collaboration. I am also grateful to Chris Barker, Gennaro Chierchia, Delia Graff, Irene Heim, Ed Keenan, Jeffrey King, Sally McConnell-Ginet, Jason Stanley, and audiences at MIT, Princeton, Northwestern, the University of British Columbia, the University of Chicago, the University of Massachusetts, Amherst, the University of Michigan, the University of Milan-Bicocca, the University of Pennsylvania, the University of Stuttgart, Yale University and the Zentrum für allgemeine Sprachwissenschaft, Berlin, for very helpful comments and feedback on earlier versions of this work. I would like to express particular gratitude to John MacFarlane and Michael Glanzburg for their thoughtful and insightful reviews of the manuscript, to Mark Richard for two challenging and extremely helpful commentaries on earlier versions of the paper, and to the Associate Editor Kent Bach for probing suggestions, questions, and comments both before and during the *L&P* review process. Any remaining shortcomings, incoherence, or vagueness are (definitely) my responsibility. This paper is based upon work supported by the National Science Foundation under Grant No. 0094263 and 0618917 and by the American Council of Learned Societies under a Charles Ryskamp Research Fellowship.

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

the basis of conventional meaning to the extent possible, allowing for context dependent truth conditions only as a last resort.

Keywords Vagueness · Gradability · Context sensitivity · Borderline cases · Sorites Paradox · Comparison class · Scale structure

1 Introduction

The general question that this paper addresses is how sentences whose main predicates consist of gradable adjectives in the (unmarked) POSITIVE FORM, such as (1), are assigned truth conditions in a context of utterance.

(1) The coffee in Rome is expensive.

The problem presented by sentences of this sort is that they are vague: what exactly it means to ‘count as’ expensive is unclear. Sentences like (1) have three distinguishing characteristics, which have been the focus of much work on vagueness in semantics and the philosophy of language. The extent to which these characteristics are interrelated is one of the questions that this paper aims to address; I begin here by presenting them from a purely descriptive perspective.

The first characteristic of vague sentences is contextual variability in truth conditions. For example, (1) could be judged true if asserted as part of a conversation about the cost of living in various Italian cities (*In Rome, even the coffee is expensive!*), but false in a discussion of the cost of living in Chicago vs. Rome (*The rents are high in Rome, but at least the coffee is not expensive!*). This kind of variability is of course not restricted to vague predicates (for example, *citizen* introduces variability because it has an implicit argument—*citizen of x*—but it is not vague), though all vague predicates appear to display it.

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

Finally, vague predicates give rise to the Sorites Paradox, illustrated in (2).

(2) *The Sorites Paradox*

- P1. A \$5 cup of coffee is expensive (for a cup of coffee).
- P2. Any cup of coffee that costs 1 cent less than an expensive one is expensive (for a cup of coffee).
- C. Therefore, any free cup of coffee is expensive.

The structure of the argument appears to be valid, and the premises appear to be true, but the conclusion is without a doubt false. Evidently, the problem lies somewhere in the inductive second premise; what is hard is figuring out exactly what goes wrong. And even if we solve this problem, we also need to explain both why it is so hard to detect the flaw in the premise, and why we are so willing to accept it as true in the first place, as pointed out by Graff (2000).

It is widely accepted that the locus of vagueness in sentences like (1) is the predicate headed by the gradable adjective *expensive*. Within linguistic semantics, a fruitful line of research has developed that analyzes the positive form as a relation between the degree to which an object possesses some gradable concept measured by the predicate and a context dependent STANDARD OF COMPARISON based on this concept. For example, *expensive* on this view denotes the property of having a degree of cost that is at least as great as some standard of comparison of cost, where the value of the standard is not part of the lexical meaning of *expensive*, but is rather determined ‘on the fly’. Truth-conditional variability arises when the standard of comparison is shifted: if the standard for *expensive* is based on the cost of coffee in Italian cities, and that is lower than the cost of coffee in Rome, (1) is true; if the standard is based on the cost of coffee in Rome vs. Chicago, and that is higher than the cost of coffee in Rome, then (1) is false.

There is general agreement among researchers in linguistics that something like this is what is going on in the interpretation of gradable predicates, but several fundamental questions remain open. The first is the question of how the semantic analysis of gradable predicates relates to an account of borderline cases and the Sorites Paradox. Specifically: what feature of the semantics of gradable predicates is responsible for their behavior with respect to these two (more general) characteristics of vague expressions? The second is the question of the actual content of the standard of comparison and how it is computed. In particular, to what extent is its value explicitly determined by the conventional meaning of gradable predicates and/or by the conventional meanings of various subconstituents of such predicates, and to what extent is it determined by purely contextual, possibly extra-linguistic factors?

These questions have received a fair amount discussion in the literature (see e.g., Barker, 2002; Bierwisch, 1989; Fine, 75; Kamp, 1975; Kennedy, 1999; Klein, 1980; Lewis, 1970; Ludlow, 1989; McConnell-Ginet, 1973; Pinkal, 1995; Sapir, 1944; von Stechow, 1984; Wheeler, 1972), but a fully comprehensive theory has not been developed. A central reason for this, I will claim, is because the full range of relevant data has not been taken into account. In particular, most analyses fail to address the distinction between RELATIVE gradable adjectives like *expensive*, which have the features of vagueness described above, and ABSOLUTE gradable adjectives like *straight* and *bent*, which do not. As I will show in detail, predicates like *straight* (as in *The rod is straight*) require their arguments to possess a maximal degree of the measured concept, and those like *bent* (as in *The rod is bent*) merely require their arguments to possess a non-zero degree of the relevant concept; neither describes a relation to a context dependent standard of comparison. However, despite these differences in interpretation, relative and absolute gradable adjectives are the same semantic type, and express the same kind of meanings. We should therefore expect an explanatorily adequate theory of the positive form to derive their differences, rather than merely stipulate them.

The goal of this paper is to develop such a theory, and in doing so, to answer the questions outlined above. I will begin by developing a semantics for the positive form of relative adjectives that improves in several ways on previous analyses, both in the way that it accounts for vagueness, and in the way it accounts for the division of labor between composition and context in the interpretation of the predicate. I will then discuss absolute adjectives and the problem they present for the analysis, providing detailed empirical evidence for a semantic distinction between relative and absolute gradable adjectives. Finally, I will propose an analysis of the relative/absolute distinction in which the truth conditions and semantic properties of the positive form—in particular, whether it is vague or not—are determined by the interaction of lexical properties of gradable adjectives, in particular the structures of the scales that represent the type of measurement they encode, and a general constraint on ‘interpretive economy’ that requires the truth conditions of a sentence based on of conventional properties of its constituents to the extent possible, allowing for context dependent truth conditions only as a last resort.

2 The semantics of the positive form

2.1 Gradable adjectives and degree morphology

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

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

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

There are various compositional implementations of the core hypotheses about gradable adjective meaning stated in (3); here I will follow Bartsch and Vennemann (1972, 1973), and Kennedy (1999) and analyze gradable adjectives as measure

¹ The leading contender to a scalar analysis of gradable adjective meaning is one that treats gradable adjectives as partial functions from individuals to truth values with context dependent extensions and adopts a supervaluational analysis of vagueness and partial models (see e.g., Fine, 1975; Ginet, 1973; Kamp, 1975; Klein, 1980; Pinkal, 1995). I will discuss these sorts of approaches in Sect. 4.4, where I will argue that even if we eventually determine that supervaluations are the best way to handle various types of semantic imprecision (including vagueness), we still need to characterize the meanings of gradable adjectives in terms of scales and degrees, since these are the features that crucially explain the relative/absolute distinction.

functions (type $\langle e, d \rangle$).² The adjective *expensive*, for example, is a function from the subset of the domain of individuals that have some cost value to (positive) degrees of cost. Measure functions are converted into properties of individuals by degree morphology, which in English includes (at least) the comparative morphemes (*more, less, as*), intensifiers (*very, quite, rather, etc.*), the sufficiency morphemes (*too, enough, so*), the question word *how*, and so forth. Degree morphemes serve two semantic functions: they introduce an individual argument for the measure function denoted by the adjective, and they impose some requirement on the degree derived by applying the adjective to its argument, typically by relating it to another degree. Syntactically, I assume that gradable adjectives project extended functional structure headed by degree morphology (Abney, 1987; Corver, 1990; Grimshaw, 1991; Kennedy, 1999), and that the adjectival projection is thus a Degree Phrase, rather than an Adjective Phrase.

As an illustration, consider the structure and interpretation of the comparative predicate in (4a), shown in (4b), where **large** represents the denotation of the adjective *large*: a measure function that maps objects to their sizes.³

- (4) a . Chicago is larger than Rome.
 b .
- DegP

$\lambda x.\mathbf{large}(x) \succ \mathbf{large}(\mathbf{Rome})$

<p>Deg'</p> <p>$\lambda y\lambda x.\mathbf{large}(x) \succ \mathbf{large}(y)$</p> <table style="width: 100%; border: none;"> <tr> <td style="text-align: center; vertical-align: top;"> <p>Deg</p> <p>$\lambda g\lambda y\lambda x.g(x) \succ g(y)$</p> <p style="margin-left: 20px;"> </p> <p style="margin-left: 20px;">more</p> </td> <td style="text-align: center; vertical-align: top;"> <p>A</p> <p>large</p> <p style="margin-left: 20px;"> </p> <p style="margin-left: 20px;">large</p> </td> </tr> </table>	<p>Deg</p> <p>$\lambda g\lambda y\lambda x.g(x) \succ g(y)$</p> <p style="margin-left: 20px;"> </p> <p style="margin-left: 20px;">more</p>	<p>A</p> <p>large</p> <p style="margin-left: 20px;"> </p> <p style="margin-left: 20px;">large</p>	<p>PP</p> <p>Rome</p> <table style="width: 100%; border: none;"> <tr> <td style="text-align: center; vertical-align: top;"> <p>than</p> </td> <td style="text-align: center; vertical-align: top;"> <p>Rome</p> </td> </tr> </table>	<p>than</p>	<p>Rome</p>
<p>Deg</p> <p>$\lambda g\lambda y\lambda x.g(x) \succ g(y)$</p> <p style="margin-left: 20px;"> </p> <p style="margin-left: 20px;">more</p>	<p>A</p> <p>large</p> <p style="margin-left: 20px;"> </p> <p style="margin-left: 20px;">large</p>				
<p>than</p>	<p>Rome</p>				

² A more standard alternative analysis is one in which gradable adjectives denote relations between degrees and individuals (type $\langle d, \langle e, t \rangle \rangle$), and comparatives and other degree constructions saturate the degree argument of the predicate; see Kennedy (1999), Heim (2000), Meier (2003), Bhatt and Pancheva (2004) and Neeleman, Van de Koot, and Doetjes (2004) for discussion of the issues at stake in choosing between the two approaches. I adopt the measure function analysis of gradable adjectives and degree morphology primarily because it provides a transparent interpretation scheme for the 'extended projection' syntax of the adjectival projection illustrated in (4b), which has a wide range of empirical and theoretical support, and because it makes it easy to see which bits of structure are contributing which bits of meaning. However, since the relational analysis also assumes that gradable adjectives encode measure functions as part of their meanings (for example, *large* holds of an individual *a* and a degree *d* just in case *a*'s size is at least as great as *d*), all of my central proposals could be adapted to this type of approach with appropriate changes in semantic type and denotation of the relevant constituents. The crucial assumptions are the ones stated in (3), which are shared by all scalar analyses of gradable predicates.

³ A note on notation: throughout this paper I follow Heim and Kratzer (1998) in assuming that syntactic representations can be directly interpreted, but I will use predicate logic as my metalanguage for representing truth conditions, rather than English as in Heim and Kratzer (1998), defining any new symbols that I introduce (as for **large** above). To keep the representations as simple as possible, I will omit type specifications for arguments from the domain of individuals, degrees and gradable adjectives, instead using the variables from the sets $\{x, y, z\}$, $\{d, d', d''\}$, $\{g, g', g''\}$, respectively. Finally, I will omit specification of assignment functions and other contextual parameters except where relevant.

Here *more* is treated as an expression that establishes an ordering relation between two degrees: one derived by applying the adjectival head to its external argument, the other by applying it to the ‘standard’ constituent, marked by *than*. Composition derives the property at the top of (4b) as the denotation of the DegP: the comparative predicate *larger than Rome* is true of an object if the degree to which it is large exceeds the degree to which Rome is large.⁴ This sort of analysis can in principle be extended to any other degree morpheme/adjective combination—collocations with *less, as, too, enough, so, how, very, quite, rather* and so forth—with appropriate modifications to the denotations of the degree morphemes.

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

Turning now to the unmarked, positive form, it is a bit paradoxical that the most morphosyntactically simple form of a gradable predicate turns out to be the hardest to adequately characterize in terms of a compositional semantic analysis.⁵ The first (and easiest) problem to overcome is a morphological/type theoretic one: if gradable adjectives denote functions of type $\langle e, d \rangle$, then they must combine with degree

⁴ I give a ‘phrasal’ semantics for *more* here for simplicity; an alternative is that the standard constituent is an elided clause that directly denotes a degree; see Hankamer (1973), Hoeksema (1984), Heim (1985), and Kennedy (1999) for discussion. I also assume that the relation between *more large* and *larger* is a matter of morphophonology.

⁵ I distinguish the completely unmarked form of e.g. *old* in (5c) from the form in (i), which is not combined with a bound degree morpheme, but is combined with a measure phrase

(i) That dog is 2 years old

The semantic justification for this distinction is twofold. First (5c) is vague and (i) is not (though it may be imprecise; see Pinkal, 1995 and the discussion in Sect. 3.2.1). Second, the two forms have quite different truth conditions: (5c) establishes a relation between the subject’s age and a contextual standard of comparison (in a way that will be described in detail below); (i) relates the subject’s age to a degree determined by the measure phrase. In terms of compositional semantics, we may assume either that measure phrases combine directly with gradable adjectives (Kennedy & McNally, 2005; Klein, 1980; Schwarzschild, 2005), or that a specialized Deg head mediates the relation between them (Svenonius & Kennedy, 2006).

morphology to derive a property of individuals. The positive form does not have any overt degree morphology, however, as illustrated by examples like (5).

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

One solution to this problem would be to take the absence of morphology at face value and assume that the positive form is a simple AP (i.e., an adjectival projection without functional degree morphology), but that the grammar includes a type-shifting rule that turns measure functions into properties of individuals (see e.g. Neeleman et al., 2004). A second solution would be to assume that the DegP in the positive form is headed by a null morpheme that has the same semantic function as overt degree morphology: it takes a gradable adjective denotation (a measure function) and returns a property of individuals (see e.g., Bartsch & Vennemann, 1972; Cresswell, 1977; Kennedy, 1999; von Stechow, 1984). In order to keep the compositional analysis of the positive form as fully parallel to the analysis of forms with overt degree morphology as possible, I will adopt the latter position, and I will refer to the relevant morpheme as *pos* (for ‘positive form’). However, nothing crucial hinges on this assumption: the content of my proposals and argumentation remains the same if we assume instead that ‘the positive degree morpheme *pos*’ is really ‘the positive type-shifting rule *pos*’.

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

In Sect. 1, I characterized the denotation of the positive form of *expensive* in terms of a relation to a contextually determined standard of comparison: *is expensive* denotes the property of having a degree of cost that is at least as great as the prevailing standard. This means that *pos* should be assigned a denotation along the lines of (6), where \mathbf{d}_s is shorthand for ‘the contextually appropriate standard of comparison, whatever that is’.

$$(6) \quad \llbracket [\text{Deg } pos] \rrbracket = \lambda g \lambda x. g(x) \succ \mathbf{d}_s$$

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

The main argument in favor of the former view is that the standard of comparison can be manipulated in what appears to be a compositional way by constituents local to the predicate, such as the *for*-PPs and modified nominals in (7a–c) and (8a–c).

- (7) a. Kyle’s car is expensive for a Honda.
 b. Nadia is tall for a gymnast.
 c. Jumbo is small for an elephant.

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

These constituents clearly have an effect on truth conditions: (7c) and (8c) could both be true in a situation in which the 'bare' positive in (9a) is false, as illustrated by the fact that (9b,c) are non-contradictory.

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

A common interpretation of these facts is that the standard of comparison is always computed relative to a *COMPARISON CLASS*, which can be made explicit by a *for-PP* or modified nominal (Klein, 1980). It may also remain implicit, however. Taking into account the fact that bare positives like (9a) can be understood as equivalent to their variants in examples like those in (7) and (8) (in appropriate contexts), many analyses assume that the comparison class is always a constituent of the semantic representation of the positive form (Bartsch & Vennemann, 1972; Cresswell, 1977; Kennedy, 1999; Kennedy & McNally, 2005; Wheeler, 1972; von Stechow, 1984).

Bartsch & Vennemann (1972), for example, provide a denotation for *pos* that is essentially equivalent to (10), where *k* is a property and **norm** is a function that returns the average degree to which the objects in the set defined by *k* (the comparison class) measure *g*.

$$(10) \quad \llbracket [\text{Deg } pos] \rrbracket = \lambda g \lambda k \in D_{(e,t)} \lambda x. g(x) \succ \mathbf{norm}(k)(g)$$

If we assume that in the absence of explicit information about the value of the comparison class argument, it can either be saturated by a variable over properties (as in e.g. Stanley, 2000) or left incomplete and 'passed up' to the propositional level (as in the variable free analysis of Jacobson, 2006; which builds on the analysis of pronouns in Jacobson, 1999), the contextual variability of the positive form boils down to the task of finding an appropriate property for the comparison class variable. This analysis therefore has the advantage resolving the context dependence of the positive form using independently necessary mechanisms for resolving property variables (which are needed to handle e.g. *one*-anaphora and other types of predicate anaphora), rather than by introducing an additional contextual parameter just for the purpose of fixing the standard of comparison of vague predicates.

Further evidence in favor of the analysis of the positive form in (10), and further support for the hypothesis that the standard of comparison is compositionally determined, comes from the fact that the standard can vary as a function of the value of the argument of the predicate. This is illustrated most clearly by an example like (11a), which can be interpreted as indicated in (11b) (Kennedy, 1999).

- (11) a. Everyone in my family is tall.
 b. for every *x* in my family, *x* has a height greater than the norm for someone like *x*, where the relevant kind of similarity (same age, same sex, etc.) is contextually determined

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

The analysis in (10) can handle this data by assuming that an implicit comparison class variable can range not just over properties, but over properties that can be relativized to bound variables, i.e., skolem functions (see Steedman, pp. 75–76 for discussion). Using \mathbf{k}_x to represent a variable over properties ‘related to x ’ (the property of being the same age as x , the property of being the same sex as x , etc.), the truth conditions of (11a) can be accurately represented as in (12).

$$(12) \quad \forall x[\mathbf{in-my-family}(x) \rightarrow \mathbf{tall}(x) \succ \mathbf{norm}(\mathbf{k}_x)(\mathbf{tall})]$$

The linguistic evidence for such functions is well-established in work on e.g. functional readings of pronouns, functional questions and choice-function analyses of indefinites, so again nothing particularly new needs to be added to the theory to account for facts like (11a).⁶

2.2 Problems with norms and comparison classes

Despite these advantages, there are a couple of problems with the analysis of the positive form outlined in the previous section. The first is that reducing the context dependence of the positive form to the identification of a comparison class fails to explain the fact that the positive form of a gradable predicate gives rise to borderline cases and the Sorites Paradox even when the comparison class is explicit (Graff, 2000; Pinkal, 1995; Rusieck, 1985). As noted above, an apparent advantage of the comparison class analysis of the positive form as implemented in (10) is that it captures context dependence in terms of general, independently motivated, interpretive mechanisms: the positive form is just another example of a construction that includes an implicit property variable. Once that variable is fixed by the contextual assignment function, the actual value of the standard of comparison can be computed strictly on the basis of the conventional meaning of *pos*: there is no additional context dependence. In particular, according to (10), the standard of comparison will always be a degree on the scale of the adjective that represents the average degree to which the objects in the comparison class possess the property measured by the adjective.

If all this is correct, however, then once the comparison class is fixed, the cutoff point for the objects that the positive form is true of is also fixed: it is the average degree to which the objects in the comparison class possess the property. As far as the truth conditions of the positive form are concerned, then, there should be no borderline cases: if an object possesses the relevant property to a degree at least as great as the average, the positive form is true of it, otherwise false. Likewise, these truth conditions should entail that the second premise of the Sorites argument is false: at some point in a Sorites sequence the average will be crossed, falsifying the

⁶ I leave aside the details of how exactly this sort of analysis would be implemented compositionally, since I will suggest an alternative below. See Ludlow (1989) and Stanley (2002) for proposals that make use of variables, and Jacobson (2006) for a variable-free analysis.

universal generalization expressed by the premise. This in itself is not a problem: if we want to retain classical logic, then the second premise must be false. The problem is that a ‘greater than average’ semantics of the positive form provides no explanation of the fact that we do not recognize the second premise to be false. (This can be seen by substituting *has a greater than average cost for expensive* in (2): the paradox disappears.)

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

Consider for example (13). The comparison class is explicit, but it is perfectly plausible that one could know that the median rent for apartments on the street is \$700 and still be unwilling to judge this sentence as true if, for example, there are a few expensive apartments with rents significantly higher than \$725.

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

Similarly, the argument in (14) remains just as paradoxical as the one discussed in Sect. 1, even though the generalization in P2 should fail for the move from \$701 to \$700.

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

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

A second response to this would be to claim that these features of vague predicates don’t have a semantic explanation in the first place, and so don’t bear on the characterization of the truth conditions of the positive in terms of an average degree (for a comparison class). For example, it could be the case that the truth conditions of the positive form are as in (10), but for purely epistemic reasons, we can never know for sure where the standard is: we know that it is an average for a comparison class, but we can never know its actual value (cf. Williamson, 1992, 1994).⁷ The

⁷ Williamson himself does not advocate a ‘greater than average’ semantics of the positive form (or any other specific semantic analysis, for that matter), though such an analysis is not inconsistent with his explanation of borderline cases and the Sorites Paradox. In fact, adopting this analysis would protect the epistemic account of vagueness from one of its central criticisms: that it is incoherent to suppose that we can know the meaning of a predicate if we can never know its actual extension. A meaning of the form ‘have a height greater than the average height in comparison class *c*’ is a perfectly coherent sense for (*pos*) *tall*, which could in principle be adduced based on generalizations over uses of the predicate, while still allowing for epistemic uncertainty about whether particular objects in the vicinity of the average (within Williamson’s ‘margin of error’) are above or below it. That said, Graff’s criticisms and the truth conditional problems discussed below remain, and the impossibility of ‘crisp judgments’ in contexts where the average and the relation of an object to it is known, discussed in Sect. 2.3 below, indicate that this account needs to be supplemented.

region of uncertainty surrounding the standard gives rise to borderline cases and the fact that we are unable to pinpoint the precise location in a graded sequence where the second premise of the Sorites Paradox fails.

However, this explanation fails to explain why we are willing to accept the second premise as true, as pointed out by Graff (2000). An even bigger problem for the specific semantic analysis in (10), observed by Bogusławski (1975), is that it does not derive the right truth conditions. If the positive form means ‘have a degree of property g greater than the average for a comparison class based on g' ’, then (15) should be a contradiction:

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

(15) is not a contradiction, however, providing clear semantic evidence that standards are not averages. This problem could be fixed by redefining **norm** so that it identifies some value other than an average for a comparison class, but any such characterization that is not itself context dependent will reproduce the problems described above. The conclusion, then, is that some element of the meaning of the positive form other than just a variable over comparison classes must be context dependent.

Before considering an analysis that implements this hypothesis, I want to address a second problem with the analysis outlined in the previous section (which is shared by many other approaches to vagueness in gradable predicates as well): the assumption that the comparison class is a semantic argument of the predicate. The evidence for this position came from *for*-PPs and modified nominals, which appear to affect the computation of the standard of comparison in a compositional way, and from the fact that they can be bound by quantifiers external to the predicate. On closer inspection, however, it turns out that none of these facts support this conclusion.

Consider first the case of modified nominals. Although such nominals typically provide the basis for computing the standard of comparison for adjectival predicates that modify them, they do not have to. As shown by (16), it is possible to assert that something is *an A NP* while denying that it is *A for an NP*.⁸

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

The conclusion to draw from (16) is that although the denotation of the modified nominal provides a highly salient property which may be used to calculate the standard of comparison, there is nothing compositional about this relation, and nothing obligatory (contra Wheeler, 1972). Consequently, there is no reason to conclude based on the tendency for a modified nominal to provide the comparison

⁸ (16) illustrates a case where the standard used by the adjective is lower than the one determined by the modified noun (i) illustrates the reverse: B’s claim that A’s assertion is false is based on a standard for the adjective that is higher than the one suggested by the noun

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

class that the latter is a constituent of the conventional meaning of the positive form, as opposed to just some other bit of relevant contextual information: the facts could be handled just as well by an analysis in which the standard of comparison is a variable whose value is determined on the basis of some discourse-salient property (cf. Pinkal, 1979). A modified noun denotation is arguably the most salient property at the point of interpreting the adjectival predicate, explaining the strong tendency for it to be used as the comparison class, but facts like (16) show that it does not have to be so used.⁹

Turning now to *for*-PPs, these appear to provide a stronger argument for positing a comparison class variable in the semantics. As illustrated by (17), it is contradictory to assert that something is *A for an NP* and simultaneously claim that it has the lowest degree on the *A*-scale for the class of *NPs*, which suggests that the *for*-PP obligatorily determines the standard of comparison.

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

Before we draw this conclusion, however, we need to consider a fact that has gone mostly unobserved in discussions of the positive form (though see Klein, 1980, pp. 13–14; Wheeler, 1972, p. 316): sentences with the structure *x is A for a NP* presuppose that *x* is an NP. For example, all of (18a–c) require Kyle's car to be a Honda, as shown by the infelicity of (19a–c).¹⁰

- (18) 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?
- (19) 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?

⁹ These observations further show that analyses that attempt to derive the interpretation of predicative uses of gradable adjectives from underlying attributive structures, such as Montague (1974) and Lewis (1970), do not actually have a theoretical advantage over analyses that treat the predicative form as basic. (Kamp, 1975 makes the same point.) Such approaches crucially assume that a modified nominal always provides a comparison class, so that structures with the form *an A NP* have non-context dependent interpretations. If predicative forms are derived from attributive ones, their context dependence can be explained in terms of principles of ellipsis, not from anything having to do with the semantics of the positive form. That is, on this view the context dependent aspect of *is expensive* involves figuring out whether it is an elided form of *is an expensive Honda*, *is an expensive Volkswagen*, *is an expensive BMW*, etc., not calculating a context dependent standard of comparison. If modified nominals do not necessarily provide comparison classes, however, then assuming a derivational relation between the predicative and attributive forms doesn't actually eliminate the problem of computing the standard of comparison in the predicative (or attributive) form.

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

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

For-PPs contrast in this regard with modified nominals, as illustrated by (20a–c).

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

These examples are somewhat odd because they don't seem to be saying much—(20a) is obviously false under normal circumstances; (20b) obviously true—but they do not involve presupposition failures. The adverbs *really*, *obviously*, etc. generate contexts that make the examples informative, but they do not affect the conclusion about the *for*-PPs: inserting them in (19a–c) effects no corresponding improvement in acceptability.

The contrast between *for*-PPs and modified nominals presents a serious challenge to the hypothesis that the comparison class is a semantic argument of the positive form, as in the analysis represented by (10). Even if we were to assume, based on (17), that a *for*-PP obligatorily determines the value of the comparison class argument introduced by *pos*, we would still fail to explain the presuppositions of an example like (19a). The interpretation we would assign to this example would be equivalent to (21), which clearly does not presuppose that Kyle's BMW is a Honda.

- (21) Kyle's BMW is as expensive as an expensive Honda.

We could augment the analysis by stipulating that the positive form presupposes that its argument must be a member of whatever set is determined by the comparison class variable, an assumption made by Klein (1980), but this would run into problems with modified nominals. As we saw above, modified nominals are not required to provide the comparison class, but they may provide it. If the argument of the adjective were necessarily presupposed to be a member of whatever comparison class is used, then we would predict that an example like (20a) should trigger a presupposition failure on an interpretation in which Kyle's BMW is asserted to (really) be a Honda whose cost is high relative to other Hondas. Such an interpretation should be blocked, but this is not the case; rather, this is the most natural interpretation of this example.

(22a,b) make the same point, in an even more striking way.

- (22) a. ??That mouse is (obviously) not small for an elephant.
 b. That mouse is (obviously) not a small elephant.

Since elephants are typically not small, the most natural interpretation of the attributive modifier in (22b) is one in which the standard is computed relative to elephants, which in turn indicates that the nominal is providing the comparison class. But (22b), unlike (22a), does not presuppose that the mouse is an elephant. This shows that it is the *for*-PP specifically, rather than the choice of comparison class more generally, that introduces the presupposition about the argument of the predicate (contra Klein, 1980).

These facts can be explained if we assume that a *for*-PP has a much more local effect on adjective meanings: it does not provide a comparison class argument for the positive form, but rather restricts the domain of the measure function denoted by the

adjective to just those objects that are members of the set defined by the nominal complement of *for* (cf. Rusiecki, 1985). That is, if *expensive* is a function from objects (that can have costs) to (positive) degrees of cost, then *expensive for a Honda* is a function from Hondas to degrees of cost, as represented in (23).¹¹

$$(23) \quad \llbracket \llbracket \text{AP expensive for a Honda} \rrbracket \rrbracket = \lambda x : \mathbf{honda}(x).\mathbf{expensive}(x)$$

According to this analysis, the examples in (19) are anomalous because the semantic argument of the measure function expressed by the AP (the subject) is not a member of its domain. The measure function fails to return a value, so the ordering relation expressed by the positive form cannot be evaluated and the sentence cannot possibly be assigned a truth value. In contrast, the examples in (20) do not run into this problem because modified nominals do not restrict the domain of the adjectives that modify them, even when they provide the comparison class.

If a *for*-PP is a modifier that restricts the domain of a gradable adjective, however, then it cannot also serve as the semantic argument of *pos*; as a result, such phrases provide no evidence for a comparison class argument. More generally, the fact that *for*-PPs and modified nominals can have the same effect on the computation of the standard of comparison in examples like (22a,b) (i.e., they ensure that it is based on the class of elephants) but different consequences for the semantic properties of the sentence (its presuppositions) show that they achieve the former result in different ways: a *for*-PP influences the standard by manipulating the domain of the adjective; a modified nominal does this by making salient the property that it denotes. A theory which assumes that both saturate a comparison class argument fails to capture this distinction.

This leaves only one piece of linguistic evidence for treating comparison classes as semantic arguments of the positive form: the ‘bound comparison class’ interpretation of examples like (24a), paraphrased in (24b).

- (24) a. Everyone in my family is tall.
 b. for every x in my family, x has a height that exceeds the standard for a comparison class based on x

In section 2.1, I claimed that the possibility of interpreting (24a) as (24b) indicates the presence of an element in the denotation of the positive form whose value can vary as a function of the denotation of the subject, which I took to be a skolemized

¹¹ I use the notational conventions for functional domain restriction from Heim and Kratzer (1998): if f is a function of type $\langle \tau, \sigma \rangle$, then $\lambda v : g(v).f(v)$ is a function just like f except that its domain is the subset of things of type τ that satisfy g . I also assume here that the NP contributes a property, though in principle it could be a set or a kind; see Graff (2000) for discussion. Note also that the property contributed by the NP part is actually not always the denotation of the NP, though it is always a function of the meaning of the NP. For example, the domain restriction contributed by the *for*-PP in (ia,b) is not the property of being a first year graduate student, but rather the property of being work done by a first year graduate student, as illustrated by the anomaly of (ic)

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

comparison class variable. However, the analysis of *for*-PPs as domain restrictors suggests an alternative explanation of this phenomenon.

If the domain of a gradable adjective can be explicitly restricted, then it is reasonable to assume that it can be implicitly restricted as well, just like the domains of other functional expressions such as quantificational determiners (von Stechow, 1994; Stanley, 2000; Stanley & Szabó, 2000b; Martí, 2002; Giannakidou, 2004; Jacobson, 2005, 2006). Of particular relevance is the fact that implicit quantifier domain restrictions can be bound in a manner fully parallel to what we see in (24), as discussed in detail by Stanley (2000). (In fact, Stanley discusses examples like (24a), and gives an analysis similar to the one I suggest in Sect. 2.1.) (25a), for example, can have the interpretation in (25b), in which the implicit domain of quantification for *exactly three* is a function of the value of the variable quantified by *most of John's classes*.

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

Let us assume that implicit domain restrictions on gradable adjectives can be bound in the same way. There are a number of different ideas about how implicit domain restrictions should be represented at the syntax-semantics interface; for the purpose of this paper, I will assume that adjectival domain restrictions are introduced via the type-shifting rule in (26), though alternative implementations are possible.¹²

¹² The approach to adjectival domain restriction outlined here differs from the typical approach to e.g. quantifier domains in locating the restriction in the (compositionally determined) presuppositions of a complex expression. In contrast, quantifier domain restrictions are usually analyzed as part of the truth-conditional meaning, specifically as conjuncts in the restrictive clause of a quantificational statement, on a par with overt (restrictive) nominal modifiers. For example, the logical form of (ia) is typically taken to be something like (ib), where \mathbf{C} is a free variable whose value is determined contextually

- (i) a. Every student passed the exam
 b. $\forall x[[\mathbf{student}(x) \wedge \mathbf{C}(x)] \rightarrow \mathbf{passed}(\mathbf{exam})(x)]$

This is not an option with gradable adjectives: conjoining a restriction to the adjective as in (iia) is impossible because it results in a type mismatch ($\mathbf{expensive}(x)$ is a degree), and conjoining it to the predicate (i.e., $\mathit{pos} A$) fails to restrict the domain of the measure function ((iib) is the meaning of *is an expensive Honda*, not *expensive for a Honda*)

- (ii) a. $*\lambda x.\mathbf{expensive}(x) \wedge \mathbf{honda}(x)$
 b. $\lambda x. \llbracket \mathit{pos} \rrbracket(\mathbf{expensive})(x) \wedge \mathbf{honda}(x)$

Interestingly, the analysis of adjectival domain restrictions presented here could be extended to the analysis of quantificational domains by generalizing the type-shifting rule in (26) (or some other implementation of the basic proposal) to noun meanings, effectively reanalyzing the restriction on a quantificational determiner's domain as a restriction on the domain of the function expressed by its nominal argument. (iii) shows the revised version of (ib); assuming \mathbf{C} is fixed in the same way in both cases, these two expressions are logically equivalent

- (iii) $\forall x[[\lambda z : \mathbf{C}(z).\mathbf{student}](x) \rightarrow \mathbf{passed}(\mathbf{exam})(x)]$.

- (26) For any gradable adjective A , there is an A' such that $\llbracket A' \rrbracket = \lambda f \lambda x$:
 $f(x) \cdot \llbracket A \rrbracket(x)$, where f is a function from individuals to truth values.

One way to saturate the domain restriction argument of the type-shifted adjective is to provide an explicit restriction: this is the job of a *for*-PP (which, as noted above, I assume denotes a property). In the absence of a *for*-PP, we may assume that this argument can be saturated either by a null proform (as in Stanley, 2000) or via compositional principles which ensure that the domain restriction is 'passed up' to the matrix (as in the variable free analysis in Jacobson, 2006).¹³ Either option will provide us with a means of ensuring that the domain restriction can vary when it is in the scope of a quantifier; this in turn means that we can derive bound comparison class readings without bound comparison classes by providing a semantics for the positive form in which the standard of comparison can be computed on the basis of the domain of the adjective. It is to such an analysis that I now turn.

2.3 A revised semantics of the positive form

The central conclusion of the previous section is that the comparison class is not the crucial feature on which we should be basing an account of the semantic properties of vague predicates. On the one hand, merely assuming that the comparison class can vary across contexts is not enough to explain the existence of borderline cases and our intuitions about the Sorites Paradox, or even the intuitive truth conditions of particular examples; we need to assume that the standard-fixing function is context sensitive as well. On the other hand, there is no clear linguistic evidence that comparison classes are actual arguments of the positive form. That is, although it is clear that a property that we can descriptively call a 'comparison class' influences the computation of the standard of comparison by providing a domain relative to which this degree is computed, this property does not correspond to a constituent of the logical form.

These considerations lead to the semantics for the positive form in (27), where s is a context-sensitive function from measure functions to degrees that returns a standard of comparison based both on properties of the adjective g (such as its domain) and on features of the context of utterance.¹⁴

¹³ In fact, Jacobson (2005) suggests that quantifier domain restrictions should be introduced in the same way as I suggest in (26) (modulo the issues raised in note (12)); if we were to adopt the 'presuppositional' analysis of nominal restrictions suggested in note (12), Jacobson's analysis would be identical to the one given here for adjectives): via a type-shifting rule that turns a nominal of type $\langle e, t \rangle$ into an expression of type $\langle \langle e, t \rangle, \langle e, t \rangle \rangle$ that is looking to combine with a restrictive modifier. If an explicit modifier is not provided, the modifier slot can be either 'passed up' to the matrix level or bound off to a quantifier using exactly the same compositional operations invoked in Jacobson's (1999) variable free analysis of (free and bound) pronouns. The result is a variable-free analysis of (free and bound) domain restrictions.

¹⁴ In form, (27) is most similar to the analyses of Lewis (1970) and Barker (2002), where s corresponds to Lewis' DELINEATION FUNCTION. The analysis differs from the Lewis/Barker analyses both the specific content ascribed to s below and in the larger role played by the lexical semantic properties of its argument, which are crucial to fixing the domain relative to which the standard is calculated, as described in detail below. At a more technical level, both the Lewis and Barker analyses treat gradable adjectives as expressions of type $\langle e, t \rangle$ (neither makes use of a semantic type 'degree', though Barker does define the truth conditions of gradable adjectives in terms of degrees), and Barker's (2002) analysis is explicitly dynamic. A dynamic component could be overlaid on the analysis developed here with no significant change to the core proposals.

$$(27) \quad \llbracket [\text{Deg} \textit{pos}] \rrbracket = \lambda g \lambda x. g(x) \succeq \mathbf{s}(g)$$

But exactly what function is \mathbf{s} ? The answer I will pursue here is that \mathbf{s} is a context-sensitive function that chooses a standard of comparison in such a way as to ensure that the objects that the positive form is true of ‘stand out’ in the context of utterance, relative to the kind of measurement that the adjective encodes.¹⁵

This proposal builds on ideas in Bogusławski 1975 and Graff 2000, but is different in an important way. For Bogusławski, the standard of comparison is the minimal degree of the measured property that is ‘conspicuous’, ‘noteworthy’ or ‘sufficient to attract attention’ for a comparison class in the context of utterance; for Graff, it corresponds to an interest-relative, ‘significant’ degree of the measured property for a comparison class in the context. Both of these characterizations imply that the interpretation of the positive form is relativized to some agent (the entity relative to whom conspicuousness, noteworthiness, or significance is assessed), but this position has recently been called into question by Stanley (2003) on the grounds that we can have beliefs about the truth or falsity of a sentence like *Mt. Everest is tall* without having beliefs about any agent relative to whom Mt. Everest’s height is supposed to be conspicuous, noteworthy or significant.

While I do not want to rule out the possibility that criteria of the sort discussed by Bogusławski and Graff play a role in some (possibly even most) cases, the notion of ‘standing out’ that I am aiming for here is designed to bypass Stanley’s objection by allowing for the possibility that the truth conditions can be satisfied based on purely distributional criteria. For example, an object in some domain can stand out relative to the kind of measurement encoded by *tall* (‘increasing height’; see Kennedy, 2001) either by having a significant or noteworthy amount of height, or by simply having a height that very few other things in the domain have, i.e. by being towards the upper end of the height continuum for the domain.

Evidence that ‘stands out’ (or something very much like it) is a matter of semantics comes from the incompatibility of the positive form with what I will call CRISP JUDGMENTS: distinctions between objects based on minor but noticeable differences in degree. For illustration of this phenomenon, consider (28a), which shows that the positive form of a gradable predicate can be used to convey the fact that there is an asymmetric ordering between two objects along some dimension, just like the (morphological) comparative form in (28b).

- (28) a. This book is long compared to that book.
 b. This book is longer than that book.

Like the comparative in (28b), (28a) does not require the positive form to actually be true of the object that it is predicated of in the context of utterance: both examples

¹⁵ An alternative would be to treat \mathbf{s} as a free variable, in which case the positive form would denote different properties in different contexts of use. This idea is the basis of contextualist analyses of vagueness (see e.g., Bosch, 1983; Kamp, 1981; Raffman, 1994, 1996; Soames, 1999), which attempt to resolve the Sorites Paradox (and explain our intuitions about the second premise) in terms of the contextual effect that the evaluation of a Sorites sequence has on vague predicate meaning—in terms of (27), on the value of \mathbf{s} . (Borderline cases receive different explanations in different accounts, depending largely on whether vague predicates are taken to denote partial functions or not.) See Stanley (2003) for a critique of this kind of analysis.

could be used to make a claim about a 100-page book in opposition to a 50-page book, for example.¹⁶ These features follow from a couple of reasonable assumptions.

First, assume that the function of the *compared to* phrase is to introduce an alternative context in which the domain of discourse includes just the two objects being compared, roughly as in (29).

- (29) $\llbracket x \text{ is } A \text{ compared to } y \rrbracket = 1$ in context c iff $\llbracket x \text{ is } A \rrbracket = 1$ in every context c' that is just like c except that the domain of discourse includes just x and y .

Second, assume a general requirement that the standard of comparison should always provide a non-trivial partitioning of the domain: one in which both its positive and negative extensions are nonempty (Klein, 1980, p. 23). It follows that in a context that includes just x and y , if it is true that x is A , it must be false that y is A . This is equivalent to saying that the degree to which x is A exceeds the standard of comparison in the *compared to*-context, but the degree to which y is A does not, which in turn entails that x is more A than y . Since the only difference between the *compared to*-context and the context of utterance is the restriction of the domain of discourse to just x and y —crucially, the degree to which x and y are A remain the same—we derive the result that a sentence of the form $x \text{ is } A \text{ compared to } y$ entails that x is more A than y .

The positive form cannot be felicitously used to express comparison in contexts involving crisp judgments, however. If instead of a 100-page book and a 50-page book, we have a 100-page book and a 99-page book, only (28b) is felicitous; we cannot use (28a) to describe such a situation. The acceptability of (28b) is unsurprising: the semantics of the comparative morphology ensures that *longer than that book* denotes a property that is true of an object just in case it has a degree of length that exceeds the length of the object denoted by *that book* (see Sect. 2.1). While slight differences in length might make it hard to evaluate the truth of such a claim, there is nothing intrinsically incompatible between the meaning expressed by the comparative form and fine differences in degree, and so nothing about the meaning of the comparative that rules out the possibility of making crisp judgments.

In contrast, given the semantics of the positive form advocated here, its meaning is inherently incompatible with crisp judgments. Given (29), an assertion of (28a) in the crisp judgment context amounts to the claim that, relative to any context including just the two books under discussion, the 100-page book is long the 99-page book is not long. But since the positive form of *long* is true of an object if and only if it stands out in length in the context of evaluation, an assertion of (28a) involves a commitment to the highly unlikely position that a difference of one page could actually be relevant to whether a book of the size of the two under consideration stands out in length or not.¹⁷ Assuming that this is a commitment that competent

¹⁶ In fact, as pointed out by Sawada (2005), the *compared to* construction actually implicates that the positive form is false of the compared objects in the context of utterance. I will abstract away from this detail in what follows.

¹⁷ One page could very well make a difference if the texts under consideration were shorter. For example, either of (ia,b) could be felicitously used to describe a 2-page essay vs. a 1-page essay

- (i) a. This essay is long compared to that one.
b. This essay is longer than that one.

speakers are not willing to make, the anomaly of (29a) follows. (I am grateful to Ede Zimmermann for suggesting this kind of characterization of the facts to me.)

The crisp judgment effect is quite general, arising not only in *compared to* constructions, but in any construction in which the positive form is used to differentiate between two objects that have only slight differences in the property measured by the adjective. In addition to collocations involving other adverbial modifiers like *relative to y*, *considering just x and y*, and so forth, we see exactly the same effects when the positive form is used in a definite description to distinguish between objects that differ only in the degree to which they possess the property measured by the adjective, as in (30a,b). (For more detailed discussion of examples like (30a,b), see Sect. 3.2.3.)

- (30) a. The long book is mine.
 b. The longer book is mine.

The definite description in (30a) can be felicitously used to refer to a 100-page book in a context in which it is paired with a 50-page book, but not in a context in which it is paired with a 99-page book; (30b) can be used in either context.

Ultimately, what is crucial to an explanation of the crisp judgment facts is that the positive form cannot be felicitously used to distinguish between two objects that differ only very slightly in some gradable property. Another way to put this is that the positive form cannot be used to distinguish between objects that are very similar relative to some gradable property, an idea that underlies several important accounts of the Sorites Paradox (in particular, Graff, 2000; Kamp, 1981; Soames, 1999; Tappenden, 1993), as well as the analysis proposed here. The explanation is essentially the same as the one proposed by Graff (2000), with one distinction.

First, starting from the semantics of the positive form in (27), we may assume that in any context, there is in fact a cutoff point between the objects that the positive form is true of and those it is false of based on the degree determined by s : the minimal degree needed to ‘stand out’ relative to g in the context. That is, we maintain bivalence for vague predicates, and assume the inductive premise of the Paradox is false. Following Williamson (1992, 1994), we may assume that the reason why we cannot say exactly which degree is the one that determines whether an object stands out relative to g is because of epistemic uncertainty about the precise location of $s(g)$ in the context; this will also account for borderline cases. Finally (and most importantly), the reason that we are unwilling to reject the second premise is the same as the reason that the positive form is infelicitous in examples involving crisp judgments: doing so would invoke a context (the one that falsifies the universal generalization) in which one of two objects that differ minimally in the property measured by g is claimed to stand out relative to g and the other not.

Like Graff (2000), I want to claim that our unwillingness to make this distinction when presented with such a pair has to do with the meaning of the positive form, though our analyses derive this result in different ways. For Graff, it has to do with the interest-relativity inherent in assessing whether an object possesses a ‘significant’ degree of some property: typically, interests in efficiency outweigh interests in discrimination of fine degrees, so two objects that differ minimally in g will always be treated the same. Since I am not positing an interest-relative semantics, I do not have access to this account. Instead, I assume with Soames (1999) that what is at work is more of a metalinguistic principle: we are unwilling to commit to the position that

one of two objects that differ minimally along the scalar continuum measured by g could stand out relative to g while the other one doesn't; as a result, whenever two such objects are evaluated together against the positive form of g , we treat them the same.

I conclude this section by showing how the semantics of the positive form proposed in (27) supports an account of the role of explicit and implicit domain restrictions on the adjective in the computation of the standard of comparison. What is crucial to the explanation is the hypothesis that the standard-fixing function is a (context-sensitive) function from the meaning of the complement of the positive degree morpheme—the adjective plus any modifiers—to a degree. Assuming for the sake of illustration that the bare adjective in (31a) has no implicit domain restriction, the measure functions relative to which the standard is determined in (31a and 31b) are not the same: one is a function from objects (that can have height) to heights, and the other is a function from gymnasts to heights.¹⁸

- (31) a. $\llbracket [\text{DegP } \textit{POS} [\text{AP } \textit{tall}]] \rrbracket = \lambda x. \mathbf{tall}(x) \succeq \mathbf{s}(\mathbf{tall})$
 b. $\llbracket [\text{DegP } \textit{POS} [\text{AP } \textit{tall} [\text{PP } \textit{for a gymnast}]]] \rrbracket = \lambda x. [\lambda y : \mathbf{gymnast}(y). \mathbf{tall}(y)](x) \succeq \mathbf{s}([\lambda y : \mathbf{gymnast}(y). \mathbf{tall}(y)])$

Given this difference in meaning between the measure functions in (31a,b), it is expected that standard of comparison picked out by \mathbf{s} in the two examples should also differ. In (31a), the standard is the degree that distinguishes objects that stand out in the context relative to a measure of height. There might be additional contextual information that restricts our attention to the class of gymnasts—for example because we are talking about gymnasts, looking at gymnasts, or using (31a) to modifying the noun *gymnast*—in which case the standard might very well be relativized to gymnasts. But this is not required. In contrast, the standard in (31b) is the degree that distinguishes those objects that stand out relative to a measure of height restricted to the class of gymnasts, i.e., to a measure of ‘height-of-a-gymnast’. In this way we derive the fact that a *for-PP* obligatorily provides the standard of comparison, though a modified nominal only optionally does so (see the discussion of (16) and (17) in Sect. 2.2).

Turning to the bound comparison class facts, if we assume that an adjectival domain can be implicitly restricted via a skolemized domain restriction variable, (32b) is a possible interpretation of (32a).

- (32) a. Every member of my family is tall.
 b. $\forall y [\mathbf{in-my-family}(y) \rightarrow [\lambda x : \mathbf{k}_y(x). \mathbf{tall}(x)](y) \succeq \mathbf{s}([\lambda x : \mathbf{k}_y(x). \mathbf{tall}(x)])]$

What varies in (32b) is not a comparison class argument, but rather the domain of the function expressed by the adjective, which is restricted to things that satisfy the contextual skolem function \mathbf{k}_y for each value of y determined by the subject. In a particular context, the value of this domain restriction might be something like ‘people that are the same age as y ’. Crucially, for each new value of y , we get a new

¹⁸ Of course, the functions denoted by *tall* and *tall for a gymnast* do not give different results (in a world at a time) for the objects that they both apply to (the gymnasts), but they nevertheless have different senses. It is this difference that is crucial here.

domain, which in turn allows for corresponding changes in the standard of comparison.¹⁹

3. Absolute gradable adjectives

3.1 Minimum and maximum standards

Most of the literature on vagueness assumes (implicitly or explicitly) that all gradable predicates in the positive form have the properties that we have observed and now analyzed for adjectives like *expensive*, *tall*, *large*, etc.: truth conditional variability, borderline cases, and Sorites sensitivity. And indeed, if the analysis presented in Sect. 2.3 is correct, then any gradable adjective is predicted to show these characteristics in the positive form. That is, if gradable adjectives are of type $\langle e, d \rangle$ and must combine with the positive form degree morpheme *pos* in order to be converted into properties of individuals, and if the properties of vagueness follow from the semantics of *pos*—that an object stand out in the context relative to the measure expressed by the adjective—then it follows that all gradable adjectives in the positive form should be vague.

In fact, this prediction does not seem to be correct. In addition to the large class of gradable adjectives that are vague in the positive form—henceforth *RELATIVE* gradable adjectives—there is a well-defined set of adjectives that are demonstrably gradable but do not have context dependent interpretations, do not give rise to borderline cases, and do not trigger the Sorites Paradox in the positive form (at least not in the way we have seen so far). Following Unger (1975), Rusiecki (1985) and more recently Kennedy and McNally (2005), I will refer to this class as *ABSOLUTE* (gradable) adjectives.²⁰

Absolute adjectives come in two varieties. *MINIMUM STANDARD* absolute adjectives, such as those in (33), simply require their arguments to possess some minimal degree of the property they describe.

- (33) a. The gold is impure.
b. The table is wet.

¹⁹ This account of variability in the standard of comparison based on quantifiers carries over directly to examples involving ‘sloppy identity’ of the standard in ellipsis constructions such as (ia) (Klein, 1980; Ludlow, 1989) and conjoined subjects as in (ib) (Cappelen & Lepore, 2005)

- (i) a. That elephant is large, and that flea is too
b. That elephant and that flea are (both) large

(ia,b) have readings in which the elephant is claimed to be large for an elephant (or more accurately, for ‘things like it’) and the flea is large for a flea. Under the assumptions outlined here, this is just another case of binding into an implicit adjectival domain restriction—directly by the subjects in (ia) and via a distributive interpretation of the conjunction in (ib)—with consequent affect on the value of the standard.

²⁰ Predicates formed out of absolute adjectives have also been studied in detail by Rotstein and Winter (2004), who refer to adjectives like those in (33) ‘partial’ predicates and adjectives like those in (34) ‘total’ predicates (using terminology introduced in Yoon, 1996). Rotstein and Winter develop semantic analyses of these predicates in terms of their underlying scalar properties, a strategy that I will also endorse in Sect. 4. They do not address the relation between relative and absolute adjectives, however, or the question of whether it is possible to develop a fully general analysis of the positive form which assigns the correct truth conditions (relative, absolute minimum, absolute maximum) given particular adjectives as inputs, which is what I am trying to do here.

- c. The door is open.
- d. The rod is bent.

Under normal usage, (33a) does not mean that the degree to which the gold is impure exceeds some context-dependent of impurity (for gold); it simply means that the gold contains some amount of impurity. Likewise, (33b) is true as long as there is some amount of water on the table; (33c) just requires some minimal positive aperture of the door; and (33d) is true of a rod that has a non-zero degree of bend.

MAXIMUM STANDARD absolute adjectives such as those in (34) require their arguments to possess a maximal degree of the property in question.

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

(34a) typically means that the platinum is totally pure, not that it merely has a lot more purity than other samples; (34b) is an assertion that the floor has no moisture on it; (34c) requires the door to be completely closed; and (34d) requires a completely straight rod.

Clear evidence that absolute gradable adjectives are gradable comes from the fact that they are perfectly acceptable in comparatives and with other degree morphology, as shown by (35)–(36).

- (35) 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.

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

They contrast in this regard with true non-gradable adjectives, which are anomalous in comparatives:²¹

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

The acceptability and interpretation of absolute adjectives in comparatives and other degree constructions indicates not only that they have the same semantic type as relative gradable adjectives—both denote functions of $\langle e, d \rangle$, and so can combine

²¹ It is often possible to coerce a gradable interpretation from a non-gradable adjective (e.g., (37b) might be understood to mean that the table has more wood in it than the door), but such interpretations are clearly marked. In contrast, the comparatives in (35)–(36) are perfectly natural and felicitous.

with degree morphology—but also that they have fundamentally the same kinds of meanings: they take an object and return a measure of the degree to which it possesses some gradable property. For example, in order to accurately capture the meaning of the comparative predicate in (36a), we simply need to assume that *pure* denotes a function **pure** from objects to degrees of purity. As illustrated in (38) (where the ordering relation expressed by the comparative morpheme *less* is the inverse of the one associated with *more*; see Sect. 2.1), this derives a meaning for the comparative that is exactly what we want: it is true of an object if it has a degree of purity that is exceeded by that of the platinum; neither object need actually be pure, or even close to it.

$$\begin{aligned}
 (38) \quad & \llbracket \text{less} \rrbracket (\llbracket \text{pure} \rrbracket) (\llbracket \text{than the platinum} \rrbracket) \\
 & = [\lambda g \lambda y \lambda x. g(x) \prec g(y)] (\mathbf{pure}) (\mathbf{the\ platinum}) \\
 & = [\lambda y \lambda x. \mathbf{pure}(x) \prec \mathbf{pure}(y)] (\mathbf{the\ platinum}) \\
 & = \lambda x. \mathbf{pure}(x) \prec \mathbf{pure}(\mathbf{the\ platinum})
 \end{aligned}$$

However, if absolute adjectives have the same semantic type and the same kind of meaning as relative adjectives, their interpretations in the positive form are unexpected. Since they are expressions of type $\langle e, d \rangle$, they must combine with degree morphology to derive a property of individuals; by hypothesis, this means combining with the degree morpheme *pos* in the positive form. But combination with *pos* should result in the sort of context-dependent interpretations of the sort discussed for relative adjectives in the previous section: *impure*, for example, should be true of an object in a context of utterance if its degree of impurity stands out, and *pure* should be true of an object if its purity stands out. But it seems that this should require something more than non-zero impurity in the former case, and something less than maximal purity in the latter, i.e. that the predicted meaning is too strong for (*pos*) *impure* and too weak for (*pos*) *pure*.

In fact, I will argue in Sect. 4 that the facts as I have described them are consistent with the analysis of the positive form given in the previous section, together with a condition that regulates the relative contributions of contextual and lexical information to truth conditions. Before moving to this discussion, however, I want to establish that the central empirical claims are correct: that absolute adjectives differ from relative ones in not giving rise to the features of vagueness that we expect given the semantics developed for relative adjectives in the previous section.

3.2 Evidence for the absolute/relative distinction

3.2.1 Imprecision vs. vagueness

The problem of absolute adjectives has not been explicitly addressed in previous work on vagueness or the semantics of gradable predicates, possibly because there is a strong initial intuition that adjectives like those in (33) actually require something more than a minimum standard, and that adjectives like those in (34) actually allow something less than a maximum standard. These intuitions are supported by examples (39a,b).

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

Awake is a minimum standard adjective, but (39a) can be felicitously uttered by someone who is not talking in his sleep. Similarly, *empty* is a maximum standard adjective, but (39c) can be used to describe a situation in which only a very few people show up to a film in a very large movie theater. These examples appear to call into question the empirical claims made in the previous section.

In fact, however, these examples illustrate a phenomenon that is distinct from vagueness, though typically exists alongside it: IMPRECISION. As discussed by Pinkal (1995), there are many expressions that have imprecise uses, but which are not vague. One clear example is predicates formed out of relative gradable adjectives and measure phrases, such as (40).

(40) The rod is 10 meters long.

(40) can be felicitously used to describe a rod whose actual length falls somewhere close to 10 meters, in a range that is itself be subject to contextual variation (e.g., 995 cm to 1005 cm in one context; 999.8 cm to 1000.2 cm in another, and so forth).

The difference between *10 meters long* and *long* is that the former, but not the latter, allows for what Pinkal (1995, pp. 99–100) refers to as NATURAL PRECISIFICATIONS: it is possible to construct a context in which *10 meters long* cleanly distinguishes between objects based on potentially very slight differences in length (e.g., a scientific experiment or a construction project); it is difficult (if not impossible) to do the same for the simple positive *long*. Put another way, it is possible to construct a natural context for the former but not the latter in which borderline cases are eliminated. This is illustrated by the examples in (41).

- (41)
- 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.

(41-a) is perfectly natural, especially if we know the context is one in which precision is important, and small differences of measurement can make a large difference in outcome (such as building a spacecraft). (41b), however, is decidedly unnatural, even in the same type of context. It is not impossible to make sense of this example, but as Pinkal points out, forcing *long* to be interpreted in this way involves fundamentally changing the meaning assigned to the positive form by the semantics.

Absolute adjectives behave like measure phrases with respect to natural precisifications. For example, assume that I am a detective in search of a violent criminal, and I'm trying to find out whether he might be hiding out in a particular movie theater. In the context described above (very few people watching a popular movie), I would consider the projectionist to be lying if he uses (39a) to respond to my question *Is anyone in the theater tonight?*. Likewise, (42) is perfectly natural in the

spacecraft construction context, even though it implies that a small amount of bend is enough to prevent the rod from counting as straight.²²

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

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

3.2.2 Entailments

Clear evidence for the relative/absolute distinction comes from entailment patterns. (The facts in this section are also discussed in Kennedy & McNally, 2005.) If the standards of comparison associated with the positive forms of absolute adjectives are minimal or maximal degrees (depending on the adjective), the truth conditions of a sentence of the form *x is A* can be characterized as in (43a) for a minimum standard adjective g_{min} , and (43b) for a maximum standard adjective g_{max} , where **min** and **max** are functions from measure functions to the minimal/maximal elements in their ranges, respectively. (Recall that I am assuming that gradable adjectives denote functions from objects to ordered sets of degrees, i.e. scales.)

²² A reviewer points out that if the amount of bend is excruciatingly small, as in (i), (42) is just as odd as (41b)

- (i) ??The rod for the antenna needs to be straight, but this one has 0.0000000001 mm bend in the middle, so unfortunately it won't work

The anomaly of (i) is presumably due to the fact that it implies that such an infinitesimally small degree of bend could actually make a difference, which is unlikely even in the construction of a spaceship. This example shows that precisifications themselves still allow for imprecision (cf. Lasersohn, 1999; Pinkal, 1995), but does not reduce the significance of the larger point here, which is simply that absolute adjectives allow for natural precisifications while relative ones do not.

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

- (43) a. $g_{min}(x) \succ \mathbf{min}(g_{min})$
 b. $g_{max}(x) = \mathbf{max}(g_{max})$

For example, plugging the measure functions expressed by *impure* and *pure* in for g_{min} and g_{max} , respectively, derives precisely the truth conditions that I described for (33a) and (34a) in Sect. 3.1: ‘*x* has greater than zero impurity’ in the former case; ‘*x* has maximal purity’ in the latter.

These meanings make predictions about the entailments of positive form absolute adjectives that are distinct from those made for positive form relative adjectives. First, (43a) predicts that a denial *x is not A_{min}* should entail that *x* possesses no amount A_{min} at all, assuming that the minimal degree on a scale represents a zero degree of the relevant property. The contradictory statements in (44) illustrate that this prediction is borne out. (# is used to indicate contradiction.)

- (44) a. #The gold is not impure, but there are some traces of lead in it.
 b. #The rod is not bent, though there is a small bend in the middle.
 c. #The door isn’t open, but it is ajar.

Second, (44b) predicts that an assertion of *x is A_{max}* should entail that *x* has a maximal amount of A_{max} , i.e., that *x* cannot be more A_{max} than it is. This sort of entailment is difficult to test, since maximum standard adjectives readily allow imprecise uses. However, as observed by Unger (1975), it is possible to force a precise interpretation by adding focal stress (specifically a falling tone) to the adjective. When we do this, as in (45), we see that the expected entailments arise:

- (45) 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 of a relative adjective in the positive form require that its argument stand out in the context relative to the measure encoded by the adjective. As a result, neither of the above entailments should hold: negation should be compatible with a positive degree of the measured property (albeit one that doesn’t stand out), and assertion should not rule out higher values. This is correct:

- (46) 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 & Winter, 2004). As shown by the examples in (47), there exist pairs of antonyms such that negation of one form entails the assertion of the other:

- (47) a. The door is not open. \Rightarrow The door is closed.
 b. The table is not wet. \Rightarrow The table is dry.
 c. The baby is not awake. \Rightarrow The baby is asleep.

The explanation for this is straightforward: both members of the pairs in (47) 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 (47) follow from the truth conditions in (43).

Relative antonyms do not show the same entailment relations:

- (48)
- a. The door is not large. \nrightarrow The door is small.
 - b. The table is not expensive. \nrightarrow The table is inexpensive.
 - c. The baby is not energetic. \nrightarrow The baby is lethargic.

Again, this follows from the fact that the standards for both positive and negative relative gradable adjectives are contextually identified and based on adjective denotation. The fact that e.g. *large* and *small* both measure size ensures that in any context, there should be some relation between their respective standards (in particular, the standard of largeness for an object should never be less than the standard of shortness for the same object), but they need not be the same degree. Something may fail to stand out with respect to the measurement expressed by *large* ('increasing size'), rendering the left-hand statement in (48a) true, and also fail to stand out with respect to the related but distinct kind of measurement expressed by *small* ('decreasing size'): it may be somewhere in the middle. This allows for the possibility that an object may be neither large nor small (which is different from being borderline).

We also see differences in the entailments of relative and absolute adjectives in comparatives. Assuming that the comparative imposes an asymmetric ordering on its arguments (see the truth conditions for the comparative in (4b) and that the truth conditions associated with minimum and maximum standard absolute adjectives in the positive form are as in (43), we predict that comparatives with absolute adjectives should generate positive and negative entailments to the positive form, respectively, depending on whether we have a minimum or maximum standard adjective (cf. Rusiecki, 1985). This prediction is borne out, as shown by the examples in (49)–(50).

- (49)
- a. The floor is wetter than the countertop. \Rightarrow
 - b. The floor is wet.
- (50)
- a. The floor is drier than the countertop. \Rightarrow
 - b. The countertop is not dry.

(49a) is true only if the floor has some degree of wetness: if it had zero wetness, then it could not possibly have a greater degree of wetness than the countertop.²³ This satisfies (43a) and generates the entailment to (49b). Similarly, in order for (50a) to be true, it must be the case that the countertop is not maximally dry (though the

²³ I assume that (49a) does not entail that the countertop also has some degree of wetness—it could in principle be completely dry (zero wetness). There is, however, a strong implicature that the countertop is also wet. This can be explained as follows: if the countertop had zero wetness, then the truth conditions of (49a) would be identical to the positive statement in (49b). The fact that the more complex comparative form is used—and in particular, the fact that the countertop is explicitly introduced as a reference point for characterizing the wetness of the floor—implicates that (49a) provides more information than the simpler positive form. But this will only be the case if the countertop has some (observable) degree of wetness itself. Note that this inference is not a pre-supposition: as shown by (i), it disappears under negation

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

floor might be). If the standard for dryness is the maximum value on the scale, as stated in (43b), 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 (51)–(52).

- (51) a. Rod A is longer than rod B. \nrightarrow
 b. Rod A/B is (not) long.
- (52) a. Rod A is shorter than rod B. \nrightarrow
 b. Rod A/B is (not) short.

This follows from the semantics of comparison and the semantics of the positive developed in Sect. 2.3: the mere fact that one object exceeds another with respect to some relative property tells us nothing about how the objects stand in relation to a contextually significant amount of the relevant property.

3.2.3 Standards of differentiation

A striking example of the context sensitivity of relative adjectives—in particular, of the way that the standard of comparison can be shifted based on salient contextual information—comes from experimental work demonstrating that a positive form relative adjective may be used as part of a definite description to distinguish between two objects that differ only in the extent to which they possess the property expressed by the adjective, even when the positive form would be judged true of neither or both objects outside of the differentiation task (see also Morreau, 2000; Sedivy, Tanenhaus, Chambers, & Carlson, 1999; Syrett, Bradley, Kennedy, & Lidz, 2005).

For example, Syrett et al. (2005) describe an experiment comparing adults' and children's ability to contextually shift a standard of comparison in which subjects are presented with examples like (53) as requests for one of two objects, which have been judged to be either both not long or both long in an independent task.²⁴

- (53) Please give me the long one.

These examples are systematically accepted as requests for the longer of the two objects, which is fully expected if the positive form of a relative adjective involves a relation to a context dependent standard of comparison. In particular, two contextual factors are at work here: the contextual focus on the two objects under consideration, and the uniqueness and existence presuppositions of the definite description. In order to satisfy the latter, the standard of comparison must be such that it makes the description true of one object but false of the other. This can be done by taking advantage of the former: if we restrict our attention to two objects that are the same except for their length, then the one that stands out relative to

²⁴ The results described here are for adults. Children behaved just like adults with minimum standard absolutes, but showed some interesting differences with maximum standard absolutes, which are analyzed in Syrett, Lidz, and Kennedy (2006a, b).

length is the longer of the two (modulo the crisp judgments effect; see the discussion in Sect. 2.3).

What is important for the current discussion is that absolute adjectives do not permit this use. (54), which contains the maximum standard adjective *full*, is accepted in a context involving a full jar and a partially full jar, but it is rejected in a context involving two partially full jars, where one is clearly fuller than the other (to a degree that controls for crisp judgments).

(54) Please give me the full one.

Similarly, (55), with the minimum standard adjective *spotted*, is accepted in a context involving a spotted disc and a spotless disc, but is rejected in a context involving a disc with just a few spots and one with many spots.

(55) Please give me the spotted one.

These results strongly support the claim that absolute adjectives have endpoint-oriented, fixed standards of comparison. The problem with (54) in the ‘non-full/less full’ context is that the existence presupposition of the definite is not met: neither jar is (maximally) full. The problem with (55) in the ‘spotted/more spotted’ context is that the uniqueness presupposition of the definite is not met: both discs are spotted. Both of these problems could be resolved, and the sentences made felicitous, if the standards of comparison for the respective adjectives could be shifted in the same way we saw for *long* above, i.e., if being (sufficiently) fuller/more spotted than another object were enough to count as full/spotted—to stand out in fullness/spot- tedness—in a context containing just those two objects. The fact that the sentences remain infelicitous, despite the pressure to satisfy the requirements of the definite, shows that the standards for *full* and *spotted* are not sensitive to the context (in the same way as relative adjectives), but are instead fixed to maximum and minimum values on the scale, respectively.²⁵

3.2.4 The Sorites Paradox

A final difference between relative and absolute adjectives involves their interaction with the Sorites Paradox. As we have seen, a characteristic of relative adjectives is that they give rise to the Paradox. This is illustrated again for the relative adjective *big* in (56).

- (56) 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.

²⁵ Note that comparative forms of both relative and absolute adjectives are felicitous in each of the conditions mentioned above (two long rods; two short rods; full/non-full container; non-full/less full container; spotted/more spotted disc; non-spotted/spotted disc)

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

This is expected: the semantics of the comparative involves a relation to an arbitrary degree—in this case, the degree to which the less long/full/spotted object possesses the property—rather than a (contextual/minimal/maximal) standard.

In Sect. 2.3, I claimed that the paradox results from semantics of the positive form, in particular from the content of the standard-fixing function s , which maps an adjective denotation to a degree that distinguishes those objects that stand out in the context relative to the kind of measure encoded by the adjective from those that don't. The reason that we fail to reject the inductive premise is because we are unwilling to treat two objects that differ slightly along the scalar continuum measured by a gradable adjective differently with respect to the property of standing out when confronted with those two objects, as in the crisp judgment context or in the falsification of the inductive premise.

If absolute adjectives in the positive form have the same kind of meaning as relative adjectives, then we might expect them to also give rise to the Sorites Paradox. This is not the case, however: (57) does not lead to a paradoxical conclusion, precisely because the second premise is quite naturally judged to be false.

- (57) 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.

(58) makes the same case for the minimum standard adjective *bent*.

- (58) P1. A rod that has 10 degrees of bend is bent.
 P2. A rod that is 1 degree less bent than a bent rod is bent.
 C. A rod that has 0 degrees of bend is bent.

The first premise is true, but the second is again easily judged false: a straight rod is 1 degree less bent than a rod with 1 degree of bend, but it is not bent.

Although it may be possible to construct contexts in which we would accept the second premises of these arguments (i.e., contexts in which precision is not important; see the discussion in Sect. 3.2.1), the important point here is that it is easy to construct contexts in which they are rejected, which is precisely what is impossible in the case of relative adjectives. Given that Sorites sensitivity is standardly presumed to be a defining characteristic of vagueness, these examples further support the conclusion that absolute adjectives are not vague. Moreover, the reason that the Paradox does not arise—the reason that we are willing to reject the inductive premises in (57) and (58)—is because the truth conditions of absolute adjectives in the positive form are as described in Sect. 3.1: *full* (as a measure of seat occupancy) means 'is maximally full', and *bent* means 'has non-zero bend'.

4 Vagueness and grammar

4.1 Natural transitions

The conclusion to be drawn from the facts discussed in the previous section—the possibility of natural precisifications, entailment patterns, the rigidity of the standard, the failure to trigger the Sorites Paradox—is that absolute gradable adjectives are not vague; instead the standard of comparison in predicates headed by absolute adjectives is conventionally associated with a minimum or maximum degree on a scale, rather than a context-dependent value, as in the case of relative adjectives.

This appears to be a problem for the analysis of the positive form presented in Sect. 2.3, as can be seen by considering the case of the maximum standard adjective *full*.²⁶ (Similar considerations apply to minimum standard absolutes; I focus on *full* for perspicuity.)

If all that were required for an object to count as full is that it stand out in the context relative to the kind of measure encoded by *full* (something like ‘extent of occupied space in a container’²⁷), then why isn’t it enough to just contain a lot of stuff relative to other objects in the context? That this is not enough is shown quite clearly by the differentiation task discussed in Sect. 3.2.3. It seems undeniable that in a context that includes just two containers with clearly different amounts of stuff inside them, the fuller of the two stands out (in some sense), yet this is evidently not enough for it to count as full if it is not in fact completely full.

Similarly, if (*pos*) *full* requires an object to stand out relative to a measure of occupied space, and if we are in general unwilling to treat adjacent objects on a gradable continuum differently with respect to this kind of property, then why don’t we see a paradox in (57)? Why don’t we resist judging the second premise false? If we replace *full* with *has a lot of occupied seats* in this argument, which is the meaning we might expect it to have given the analysis proposed in Sect. 2.3, the paradox returns (assuming the theater has a decent number of seats in the first place), which shows that this is not what (*pos*) *full* actually means.

However, a closer look at our intuitions about absolute adjectives in the Sorites Paradox suggests a way of reconciling their endpoint-oriented truth conditions in the positive form with the ‘stands out’ semantic analysis developed for relative adjectives in Sect. 2.3. Recall that the explanation of the Sorites Paradox presented there (specifically of our unwillingness to reject the inductive premise) relied on the idea that two objects that differ very slightly on a graded continuum are judged the same relative to the positive form because treating them differently would involve a commitment to the unlikely position that such a small difference in degree could matter as to whether one stands out or not, relative to the measure on which the continuum is based.

Implicit in this hypothesis is the idea that there is nothing about the two adjacent degrees onto which the objects are mapped that distinguishes them from each other, or from any other pair of adjacent degrees on the scale (other than relative ordering). To use a term from Williamson (1992), there is no ‘natural transition’ on the basis of which we would be willing to say that one object stands out but the other

²⁶ At this point it might appear that an analysis in which vagueness is a (lexical) property of adjective meaning rather than the positive form (as in Barker, 2002; Fine, 1975; Kamp, 1975; Klein, 1980; Lewis, 1970) would be better equipped to handle the data, since it could be captured as a lexical distinction. This conclusion would be incorrect, however. Such an account would explain the difference between relative and absolute adjectives by analyzing the former as context-sensitive, partial functions from individuals to truth values, and the latter as complete, non-context dependent functions from individuals to truth values. But this would incorrectly predict that absolute adjectives should be unacceptable in comparatives, since the semantics of comparison in such analyses is built on top of the analysis of vagueness, typically by quantifying over the different ways that a vague predicate could be made precise. In this type of approach, then, vagueness is a necessary condition for gradability; it follows from the fact that absolute adjectives are gradable but not vague that the absolute/relative distinction cannot be explained in these terms. I come back to this issue in more detail in Sect. 4.4.

²⁷ Note that there is a kind of indeterminacy in *full* (and *empty*) as to whether degree of occupied space is based on total physical space or conventionalized units, such as seats in a theater, tables in a restaurant, etc.

doesn't. The situation is different with absolute adjectives, however. Intuitively, the reason that we are willing to judge the inductive premises false in (57) and (58) is because the kinds of measurement that the adjectives encode come with natural transitions relative to which we can distinguish two objects: the move from a non-maximal degree to a maximal one in the case of *full*, and the move from a minimal degree (i.e., zero) to a non-minimal one in the case of *bent*.

Given these observations, we may hypothesize that what it means to stand out relative to the measure expressed by an absolute adjective is simply to be on the upper end of one of these transitions, i.e., to have a maximal or non-zero degree of some property as opposed to a non-maximal or zero degree (depending on the adjective; more on this in the next section). This would derive the desired truth conditions for absolute adjectives, providing an explanation of the facts discussed in Sect. 3, and would allow for a unified account of the truth conditions and semantic properties of both relative and absolute adjectives in the positive form—in particular, the fact that the former are vague but the latter are not—in terms of the analysis presented in Sect. 2.3.

However, in order for this analysis to have any explanatory power, we need to say first why only absolute adjectives allow for 'natural transitions', and second why such transitions *must* be used to fix the standard of comparison: if doing were merely an option, and not required, then we would expect an adjective like *full* to be ambiguous between an absolute and a relative interpretation, contrary to fact. In the following sections, I show that these results follow from two factors: the structure of the scale used by an adjective and a general principle of interpretive economy, which requires truth conditions to be based on the conventional meanings of the elements of a sentence to the extent possible, allowing for context sensitivity only as a last resort.

4.2 Scale structure and standard of comparison

The hypothesis that gradable adjectives denote functions from individuals to degrees (or that they incorporate such functions as part of their meanings; see note 2) leads to the expectation that gradable adjectives may differ with respect to features of the sets of degrees that comprise their ranges, i.e., their scales. There are a number of different ways that adjectival scales can be formalized, but minimally they must be triples $\langle D, \prec, \delta \rangle$ where D is a set of points, \prec is a total ordering on D , and δ is a dimension (see Bartsch & Vennemann, 1972, 1973; Bierwisch, 1989; Kennedy, 1999).

The dimension indicates the kind of measurement that the scale represents, and is the most obvious parameter of scalar variation, since it both distinguishes different adjectives from each other (e.g. *expensive* measures an object along a dimension of cost while *fast* measures an object along a dimension of speed), and different senses of individual adjectives from each other (e.g., *long* can measure an object either with respect to linear extent or temporal extent; see the discussion of indeterminacy at the end of Sect. 2.1).

The scalar feature that is relevant to the current discussion is the structure of the set D , which provides the range of the function expressed by the adjective: the set of degrees. Kennedy and McNally (2005, argue that gradable adjectives may differ with respect to the structure of this set, in particular, whether it does or does not have minimal or maximal elements, i.e., whether the scale is open or closed (see also

Kennedy & McNally, 1999; Paradis, 2001; Rotstein & Winter, 2004). There are four obvious ways in which scales could vary according to this feature, which are schematically represented in (59): a scale could lack either minimal or maximal elements, it could include a minimum but no maximum, it could include a maximum but no minimum, or it could include both a minimum and a maximum.²⁸

- (59) *A typology of scale structures*
- a . (TOTALLY) OPEN: ○————○
 - b . LOWER CLOSED: ●————○
 - c . UPPER CLOSED: ○————●
 - d . (TOTALLY) CLOSED: ●————●

As shown independently by Rotstein and Winter (2004) and Kennedy and McNally (2005), evidence for this typology comes from the distribution of adjectival modifiers that pick out maximal and minimal degrees on a scale. The reasoning relies on the following observation (which goes back at least to Sapir, 1944): antonymous pairs of gradable adjectives map their arguments onto the same scale (*tall* and *short* both measure degree of height, *wet* and *dry* both measure amount of moisture, *full* and *empty* both measure level of contents, and so forth), but impose inverse orderings on their shared domains. This is reflected by equivalences like (60).

- (60) *a* is {taller, fuller, wetter} than *b* ⇔
b is {shorter, emptier, dryer} than *a*.

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

There are different ways to formally capture this relation between polar antonyms, which depend on particular assumptions about the representation of degrees.²⁹ What is important for the current discussion is that the relation between polar antonyms and their shared scales leads to the following predictions if the scalar typology in (59) is real. First, if a scale is closed on the lower end, then the range of the positive member of an antonym pair that uses that scale should

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

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

include a minimum degree (one that is ordered below all others) and the range of the negative member should include a maximum degree (one that is ordered above all others). Second, if a scale is closed on the upper end, then the range of the positive member of an antonym pair should include a maximum degree and the range of the negative member should include a minimum degree. Conversely, if a scale is open on the lower end, then the positive antonym should have no minimum degree and the negative antonym should have no maximum; and if a scale is open on the upper end, the positive should have no maximum and the negative no minimum.

The empirical probe that Rotstein and Winter (2004) and Kennedy and McNally (2005) use to test these predictions is acceptability of degree modifiers that pick out maximal or minimal degrees on the scales of the adjectives they modify. *Absolutely*, *completely*, *totally* and *perfectly* are examples of the former kind of degree modifier; *slightly* and *partially* are examples of the latter type. The general set of predictions for combinations of these modifiers with positive and negative adjectives using the four scale types illustrated above are laid out in (61).

(61)	OPEN	L-CLOSED	U-CLOSED	CLOSED
	Deg _{max/min}	Deg _{max/min}	Deg _{max/min}	Deg _{max/min}
A _{pos}	??/??	??/√	√/??	√/√
A _{neg}	??/??	√/??	??/√	√/√

The empirical picture is complicated somewhat by the fact that not all modifiers co-occur with all adjectives for apparently idiosyncratic reasons, and some of the maximizers can also be assigned ‘high degree’ rather than strictly maximal interpretations (see Kennedy & McNally, 2005 for discussion of how to control for this). However, the modifiers *perfectly* and *slightly* provide clear judgments across a broad number of cases, and as shown by the examples in (62), demonstrate that the expected pattern does in fact emerge.

- (62) *Open scales*
 a. ??perfectly/??slightly {tall, deep, expensive, likely}
 b. ??perfectly/??slightly {short, shallow, inexpensive, unlikely}
- (63) *Lower closed scales*
 a. ??perfectly/slightly {bent, bumpy, dirty, worried}
 b. perfectly/??slightly {straight, flat, clean, unworried}
- (64) *Upper closed scales*
 a. perfectly/??slightly {certain, safe, pure, accurate}
 b. ??perfectly/slightly {uncertain, dangerous, impure, inaccurate}
- (65) *Closed scales*
 a. perfectly/slightly {full, open, opaque}
 b. perfectly/slightly {empty, closed, transparent}

There are a number of questions that should be asked about these patterns, most important of which is: why do the scales used by particular adjectives have the structure they do? For example, naive intuition suggests that the *COST* scale should

have a minimal value representing complete lack of cost, just as the DIRT scale has a minimal value representing complete lack of dirt. However, the unacceptability of *??slightly??partially expensive* and *??perfectly??completely??absolutely inexpensive* (cf. *slightly/partially dirty* and *perfectly/completely/absolutely clean*) indicates that as far as the gradable adjective pair *expensive/inexpensive* is concerned, this is not the case: the scale used by these adjectives to represent measures of cost does not have a minimal element. Instead, the property of zero cost is named by the non-gradable adjective *free*. The structure of a scale is presumably determined mainly by the nature of the property that it is used to measure, but the different behavior of e.g. *expensive/inexpensive* vs. *dirty/clean* suggests that this aspect of linguistic representation may diverge from what naive intuitions suggest.

Whatever principles explain why particular adjectival scales have the structures they do (see Kennedy & McNally, 2005 for specific proposals regarding the scale structure of adjectives derived from verbs), what is important for the current discussion is that the scalar typology illustrated in (59) has direct consequences for the interpretation of the positive form. As observed by Kennedy and McNally (2005) and as illustrated by the examples in (62)–(65), gradable adjectives that use totally open scales have relative interpretations in the positive form; gradable adjectives that use totally or partially closed scales have absolute interpretations. (Rotstein and Winter (2004) also observe the correlation between closed scales and absolute interpretations, but do not discuss the open scale/relative adjective correlation.)

The first generalization is exceptionless: open scales by definition lack maximal and minimal values, and so the endpoint oriented standards associated with absolute adjectives are simply unavailable. More generally, there is nothing inherent to the structure of an open scale that results in natural transitions: open scales represent infinitely increasing or decreasing measures. As a result, there is nothing about the meaning of an open scale adjective alone that provides a basis for determining whether an object stands out relative to the kind of measure it encodes. In order to make such a judgment, we need to invoke distributions over domains relative to which a standard can be established; i.e., we need a comparison class, which can be provided either by the context or by the adjective via a domain restriction, as we have seen. Even when we restrict our attention to a comparison class, however, the notion of standing out relative to an open scale measure remains inherently unstable and subject to the kinds of considerations discussed in Sect. 2.3 that give rise to the vagueness of relative adjectives.

The second generalization is not quite exceptionless, but the facts discussed in Sect. 3 clearly show that in general, the standard of comparison of adjectives whose scales are closed on one end or the other corresponds to an endpoint of the scale: the minimum in some cases (*impure, dirty, bent*, etc.); the maximum in others (*pure, clean, straight*, etc.).³⁰ Refining my suggestions above, these facts can be explained in terms of the semantics of the positive form developed in Sect. 2.3 if we assume that what it means to stand out relative to the measure expressed by a closed scale adjective is to be on the upper end of a natural transition based on the scale: the transition from a

³⁰ Perhaps the most notable exception to this generalization is the philosophers' favorite vague adjective, *bald*. Modification patterns indicate that it has a closed scale (*almost/completely/perfectly bald*), but relative to the tests in Sect. 3.2, it behaves like a relative adjective. This is the exception that proves the rule, however: the fact that the majority of closed scale adjectives are not like *bald* shows that there is something to explain here.

non-maximal degree to a maximal one or the transition from a zero degree to a non-zero degree. Crucially these transitions are inherent to a closed scale and to the kind of measure it represents, and so also to the meaning of a closed scale adjective. In short: a maximal degree always stands out on an upper closed scale, and a non-minimal degree always stands out on a scale with a minimum (zero) element. (I will address the question of scales with *both* maximum and minimum elements below.)

This can be only part of the story, however. Having a closed scale is a sufficient condition for absolute truth conditions, but not a necessary one. That is, even if we accept the reasoning just articulated, we don't rule out the possibility that something less than a maximal degree would be enough to stand out, or that something more than a (merely) non-minimal one would be required. That is, there is no inherent semantic incompatibility between (totally or partially) closed scales and a relative interpretation of the positive form, so we might expect that relative truth conditions are always an option for absolute adjectives, contrary to fact.

4.3 Interpretive economy

I propose to account for the default mapping of closed scale adjectives to absolute interpretations in terms of a principle of *Interpretive Economy*, which requires truth conditions to be computed on the basis of the conventional meanings of the expressions of a sentence (or logical form) to the extent possible, allowing for context-dependent truth conditions only when conventional meaning is insufficient. This constraint is stated in the form of the directive in (66).³¹

(66) *Interpretive Economy*

Maximize the contribution of the conventional meanings of the elements of a sentence to the computation of its truth conditions.

The intuition that *Interpretive Economy* is designed to capture is that although participants in a discourse may not be in full agreement about those properties of the context that play a role in the computation of context-dependent features of meaning, they are in agreement about the conventional meanings of the words and complex expressions in the sentences they use to communicate (assuming they share the same lexicons and grammars). (66) requires maximization of the role of agreed upon meanings, simplifying composition and constraining the range of possible interpretations that a particular structure may be assigned.

The effect of *Interpretive Economy* on the positive form is to ensure that closed scale adjectives are absolute. According to the reasoning in the previous section, the natural transitions at the edges of a closed scale provide a basis for fixing the standard of comparison in a way that is consistent with the meaning of the positive form (that an object that satisfies it stands out relative to the measure the adjective expresses). A context-dependent, relative standard of comparison is also in principle an option, but since an adjective's scale structure is part of its conventional meaning,

³¹ An important question is what kind of constraint (66) is. For the purpose of this paper, I will assume that it is a constraint on semantic processing. This is clearly a hypothesis that should be further tested and developed in an experimental context; here I will focus on providing empirical support for it by showing the role it plays in explaining the facts discussed in this paper.

Interpretive Economy dictates that the absolute truth conditions are the ones that should surface.

For illustration, consider the case of *bent/straight*, which use a lower-closed scale according to the diagnostics discussed in the previous section. For the positive adjective *bent*, both a minimum standard interpretation and relative one are in principle possible: the standard could either correspond a degree of bend that distinguishes objects based on whether they stand out relative to the natural transition provided by the lower closed scale (the move from zero to positive bend) or a degree of bend that distinguishes those objects in some comparison class that stand out in degree of bend from the others. Since the former truth conditions are based strictly on the lexical meaning of the adjective (its scalar properties), *Interpretive Economy* selects this meaning. The analysis of *straight* is essentially the same, except that the choice is between a relative and a maximum standard interpretation. Again, the absolute (maximum) interpretation wins out because it involves truth conditions that are based strictly on the meaning of the adjective: the fact that it uses a scale that has a maximum value (complete straightness/zero bend).

Interpretive Economy makes clear predictions about gradable adjectives with lower and upper closed scales, since the only competing truth conditions are those associated with the relative and absolute interpretations (whichever is appropriate given the scale structure of the adjective). Adjectives with totally closed scales are a bit more complicated, but ultimately provide important confirmation of the hypothesis that there is a general principle like (66) at work here. All other things being equal, totally closed scales are compatible with each of the three kinds of interpretations we have seen: relative, minimum absolute, and maximum absolute. *Interpretive Economy* rules out the relative interpretation, but on its own it says nothing about the choice between the two absolute ones. If such a principle is at work, then, the prediction is that closed scale adjectives should display interpretive variability in the positive form, taking on maximum standard interpretations in some contexts and minimum standard interpretations in others.

The antonyms *opaque* and *transparent* verify this prediction. According to the diagnostics discussed above, these use a totally closed scale (*completely/slightly opaque/transparent*), and so are in principle compatible with either minimum or maximum standard interpretations in the positive form. The following examples show that both interpretations are in principle possible. Consider a context in which I am manipulating a device that changes the degree of tint of a car window from 0% (completely transparent) to 100% (completely opaque). (67a) can be felicitously uttered 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).

- (67) 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, (67b) 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).

Out of context, there is a preference for maximum standard interpretations of *opaque* and *transparent*, but this can be explained in pragmatic terms: for any closed

scale adjective, a maximum standard interpretation entails a minimum one, but not vice-versa. Assuming that stronger meanings are in general favored (cf. Dalrymple, Kanazawa, Kim, Mchombo, and Peters' (1998) analysis of the interpretive variability of reciprocal constructions), this preference follows. If minimum standard interpretations were impossible, however, then the second sentences in (67a,b) would be contradictory.

Other total scale adjectives that allow both minimum and maximum standard interpretations are *open*, *exposed* and *uncovered*. The examples in (68) illustrate this for *open*.³²

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

If I am a member of the crew of the starship *Enterprise* and I do not understand (68a) to be a warning that any amount of opening of the airlock will result in depressurization, then I am a danger to the ship and crew. Likewise, if I am the helmsman and fail to understand (68b) as an prohibition against trying to leave the station before the space door is completely open (here *the space door* refers to the door of the space station, which the ship needs to pass through in order to get into space), I am again a danger to the ship and crew.

Despite these examples, however, many totally closed scale adjectives have fixed standards of comparison. For example, *closed*, *hidden* and *covered*, the antonyms of *open*, *exposed* and *uncovered*, all have maximum standards; so do the pair *full/empty*. Others, such as the deverbal adjectives *acquainted* (*with y*) and *documented* have minimum standards. (These claims can be verified by examining the adjectives' entailment patterns; see Sect. 3.2.2.) We could say that these adjectives are lexically specified for particular absolute standards, but this would undermine the hypothesis that there is a general principle like *Interpretive Economy* at work here: if we have to invoke lexical specifications for some absolute adjectives, there would be little reason to think that all of them are not handled in this way.

In fact, closer inspection of a wide array of totally closed scale adjectives reveals a systematic relation between features of their meanings and the orientation of their absolute standards, as we would expect given (66), but not if the orientation of the standard is arbitrarily specified for particular adjectives. The picture is clearest if we focus on the case of deverbal adjectives. As pointed out by Kennedy and McNally (2005), there is a correlation between the scale structure of a deverbal adjective and the event structure of its source verb: if the verb is telic, its adjectival form has a totally closed scale; if it is atelic, it has a partially closed scale. Kennedy and McNally further show that in the specific case of total scale deverbal adjectives derived from accomplishment verbs, the orientation of the absolute standard (minimum vs. maximum) is a function of the role played by the argument of the adjective in the event described by the verb. If the argument corresponds to an incremental theme, then the adjectival form has a maximum standard; if it does not, then the adjectival form has a minimum standard.

³² Similar examples can be constructed for *exposed* and *uncovered*. For example, an exposed line of troops is a line of troops in which some soldiers are exposed to enemy fire (as pointed out to me by Mark Richard). However, the uproar over Janet Jackson's exposed breast during the halftime show of Superbowl XXXVIII resulted not because there was some visible skin, but because her breast was completely exposed, i.e., exposed to the nipple.

This difference is illustrated quite clearly by the adjectival form(s) of *loaded*. When the verb is turned into an adjective that measures the degree to which its goal argument is loaded with some contents (*loaded with*), it has a minimum standard; when it is turned into an adjective that measures the degree to which its incremental theme argument is loaded into a container (*loaded on*), it has a maximum standard. (See Dowty, 1991 for arguments that the ‘contents’ argument of *load* is the incremental theme.) This is illustrated by the contrast in (69): (69a) is consistent with a non-maximal degree of ‘loaded with-ness’; (69b) is inconsistent with a non-maximal degree of ‘loaded on-ness’.

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

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

These predictions are verified by further examination of *loaded*. (69b) showed that it has only a maximum standard interpretation with an incremental theme argument. (If it also allowed a minimum standard interpretation, then (69b) would not be a contradiction; cf. *The boxes are partially loaded on the truck.*) (69a) showed that *loaded* can have a minimum standard when predicated of a non-incremental theme argument; (70) shows that it can also be assigned a maximum standard.

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

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

Most adjectives with totally closed scales appear to be deverbal (perhaps because such scales are so closely tied to event structures, as argued by Kennedy and McNally), though not all are related to accomplishment verbs. The pairs *open/closed*, *covered/uncovered*, *hidden/exposed* are based on achievements, but the core of the explanation outlined above carries over directly. The difference between the members of such adjectival pairs is that the verbal source of one member (*open*, *uncover*, *expose*) names the initiation

point of an event, while the verbal source of the other (*close, cover, hide*) names the culmination point. As a result, an argument of the ‘initiation predicate’ will count as having participated in the event named by the verb just in case it has a minimal degree of the measured property, while an argument of the ‘culmination predicate’ will count as having participated in the event just in case it has a maximal degree of the relevant property. The latter are therefore predicted to have maximal standards, and the former to have variable standards, as we have seen.

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

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

Assuming the stage/individual contrast correlates with the presence of an event variable in the predicate, this contrast suggests that *full/empty* are eventive in a way that *opaque/transparent* are not. Their interpretations in the positive form then follow from the considerations outlined above. The positive form of *full/empty* introduces an entailment that their arguments have (maximally) participated in an event of filling/emptying; this rules out a minimum standard interpretation. The positive form of *opaque/transparent* has no such entailment, so both the maximum and minimum standard interpretations are in principle possible, though pragmatic considerations should favor the former.

4.4 Vagueness and gradability

The assumptions about gradable adjective meaning that I adopted at the beginning of this paper, in which gradable adjectives encode measure functions that map objects to degrees on a scale, which in turn have crucial representational properties (the open/closed distinction), constitute one of two main approaches to the semantics of gradable predicates. The other kind of approach, taken in supervaluation analyses of vagueness and developed in most detail as a compositional theory of gradable predicate meaning by Klein (1980) (see also Fine, 1975; Larson, 1988; Kamp, 1975; McConnell-Ginet, 1973), does not make use of degrees (either in the type system or in the semantic ontology), but instead analyzes gradable predicates in a way parallel to other predicative expressions, as functions of type $\langle e, t \rangle$. What is special about gradable predicates is that they denote context-sensitive, partial functions from individuals to truth values: in addition to their positive and negative extensions, they may have an ‘extension gap’, which corresponds to the set of objects that the predicate is neither true nor false of in a particular context of utterance.

Crucially, the positive and negative extensions and the extension gap of a gradable predicate may vary across contexts of use, becoming more or less precise. This variability underlies the properties of vagueness that have been the focus of this paper, and provides the basis for a semantics of comparatives and other complex forms.³³ Comparatives, for example, involve quantification over possible interpretations (or ‘precisifications’) of an adjective: *x is more A than y* is true just in case there is an interpretation of *A* such that *x* is in its positive extension and *y* is not. In effect, this type of approach derives gradability from a general semantics for vague predicates, while degree-based approaches build an account of vagueness on top of a more general semantics of gradability.

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

In this paper, I have argued that the explanation for this difference is based on the structures of the scales onto which relative and absolute gradable adjectives map their arguments. In a degree-based semantics of gradable adjectives, scale structure is a basic component of the conventional meaning of a gradable adjective—it is a feature of the range of the measure function the adjective encodes—and as such, it may be accessed in the computation of complex meanings. Rotstein and Winter (2004) and Kennedy and McNally (2005) show how scale structure interacts with the semantics of various kinds of degree morphemes; here I have claimed that it interacts with the principle of *Interpretive Economy*, effectively restricting the kinds

³³ I should note that the contextual variability of the function expressed by the base adjective must also be subject to general constraints that reproduce the effects of a scalar semantics. For example, if an object *x* is in the positive extension of *A* in context *c*, and $x \succ y$ with respect to the property described by *A*, then there is no context *c'* such that *y* is in the positive extension of *A* in *c'* but *x* is not; see e.g. Klein’s (1980) *Consistency Postulate* and Fine’s (1975) ‘penumbral connections’

Of course, a degree-based semantics of gradable adjectives is also compatible with this approach to vagueness. That is, we could assume both that gradable adjectives are type $\langle e, d \rangle$ and combine with a special degree morpheme to generate a property of individuals, as we have in this paper, and that the resulting property denotes a partial function from individuals to truth values that has the features that these kinds of analyses postulate. This would derive the consistency postulate, and would provide a way of circumventing the criticisms I articulate below, but would also be subject to the various arguments that have been leveled against supervaluation-based analyses of vagueness (see e.g. Graff, 2000; Williamson, 1992, 1994).

of truth conditions that a closed-scale adjective can have in the positive form. The vagueness of the positive form is thus a function both of compositional semantics (via the meaning of *pos*) and lexical semantics (the scale structure of a gradable adjective), mediated by *Interpretive Economy*. It is in this sense that vagueness is directly influenced by the grammatical properties of the expressions that (potentially) give rise to it.

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

5 Conclusion

This paper has used the semantic properties of relative and absolute gradable adjectives in the positive form as an empirical probe on the relation between vagueness and grammar, arguing that whether the positive form of a gradable adjective is vague is a function of two factors: the structure of the scale that represents the kind of measure that the adjective encodes, and a general principle of *Interpretive Economy* which maximizes the role of conventional meaning to the calculation of truth conditions, allowing for the contribution of contextual information only as a last resort. The specific analysis of vagueness that I have advocated is one that is based on the semantics of the positive form, which requires an object to stand out relative to the measure encoded by the adjective. In the case of adjectives with closed scales, the conditions under which this property holds are based strictly on properties of the scale, resulting in either an absolute minimum or absolute maximum interpretation. In the case of adjectives with open scales, it is necessary to look to the distribution of objects in some domain (a comparison class) relative to the measure function the adjective expresses in order to fix truth conditions for the predicate; vagueness comes from epistemic uncertainty about where we actually draw the line and metalinguistic resistance to treating highly similar objects differently relative to the property expressed by the positive form.

Whether this analysis extends to an account of vagueness in other categories is an issue that must be addressed in future work. While it should be straightforward to apply the proposals here to the case of vague determiners like *many/few* and *much/little* (construed as measure functions on pluralities and substances, respectively; see Hackl, 2000) and scalar nouns like *heap, pile* and so forth, it is less clear that the analysis will extend to category-denoting nouns like *tree* or *planet*—or that it should.

Such terms can have borderline cases, but they are less easily associated with Sorites sequences, and can be given natural precisifications.³⁴ This suggests that there is more than one kind of vagueness (Soames, 1999), or perhaps that there are different kinds of ‘interpretive indeterminacy’ that have some shared properties. I believe that the resolution of these issues depends to a large extent on a better understanding of a second important question left open in the current study: the relation between vagueness and imprecision. These notions are often conflated, either by characterizing vagueness as a kind of imprecision (e.g. Lewis, 1970, 1979; Pinkal, 1995) or by treating all imprecision as vagueness (Austin, 1961, Travis, 1996). If the arguments in this paper are correct, then these two kinds of interpretive indeterminacy are distinct: absolute adjectives are not vague, but they allow for imprecision. If imprecision is an inherently pragmatic phenomenon (Lasersohn, 1999), this suggests that vagueness is a matter of semantics, arising only when the conventional meanings of particular constituents conspire to produce it, while imprecision is a more general matter of use.

References

- Abney, S. P. (1987). The English noun phrase in its sentential aspect. Ph.D. thesis. Cambridge, Mass: Massachusetts Institute of Technology.
- Austin, J. L. (1961). The meaning of a word. In J. O. Urmson, & G. J. Warnock (Eds.), *Philosophical papers* (pp. 55–75). Oxford: Oxford University Press.
- Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25(1), 1–36.
- Bartsch, R., & Vennemann, T. (1972). The grammar of relative adjectives and comparison. *Linguistische Berichte*, 20, 19–32.
- Bartsch, R., & Vennemann, T. (1973). *Semantic structures: A study in the relation between syntax and semantics*. Frankfurt: Athänum Verlag.
- Bhatt, R., & Pancheva, R. (2004). Late merger of degree clauses. *Linguistic Inquiry*, 35, 1–46.
- Bierwisch, M. (1989). The Semantics of gradation. In M. Bierwisch, & E. Lang (Eds.), *Dimensional adjectives*. (pp. 71–261). Berlin: Springer-Verlag.
- Bogusławski, A. (1975). Measures are measures: In defence of the diversity of comparatives and positives. *Linguistische Berichte*, 36, 1–9.
- Bosch, P. (1983). “Vagueness” is context-dependence. A solution to the sorites paradox. In T. Ballmer & M. Pinkal (Eds.), *Approaching vagueness* (pp. 189–210). Amsterdam: North Holland.
- Cappelen, H., & Lepore, E. (2005). A tall tale: In defence of semantic minimalism and speech act pluralism. In G. Preyer, & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning and truth*. Oxford: Oxford University Press.
- Corver, N. (1990). The syntax of left branch extractions. Ph.D. thesis, Tilburg University.
- Cresswell, M. J. (1977). The semantics of degree. In B. Partee (Ed.), *Montague grammar* (pp. 261–292). New York: Academic Press.
- Cruse, D. A. (1986). *Lexical semantics*. Cambridge, UK: Cambridge University Press.
- Dalrymple, M., Kanazawa, M., Kim, Y., Mchombo, S., & Peters, S. (1998). Reciprocal expressions and the concept of reciprocity. *Linguistics and Philosophy*, 21(2), 159–210.
- Dowty, D. R. (1991). Thematic proto-roles and argument selection. *Language*, 67, 547–619.
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30, 265–300.

³⁴ At the time of completion of this paper, the International Astronomical Union was voting on a proposed definition of *planet* which reads as follows: “A planet is a celestial body that (a) has sufficient mass for its self-gravity to overcome rigid body forces so that it assumes a hydrostatic equilibrium shape, and (b) is in orbit around a star, and is neither a star nor a satellite of a planet.” Whether an object is a satellite or not is determined by the following addendum: “For two or more objects comprising a multiple object system, the primary object is designated a planet if it independently satisfies the conditions above. A secondary object satisfying these conditions is also designated a planet if the system barycentre resides outside the primary. Secondary objects not satisfying these criteria are ‘satellites’. Under this definition, Pluto’s companion Charon is a planet, making Pluto-Charon a double planet”.

- Fintel, K. von (1994). Restrictions on quantifier domains. Ph.D. thesis, University of Massachusetts, Amherst.
- Giannakidou, A. (2004). Domain restriction and the arguments of quantificational determiners. In R. Young (Ed.), *Proceedings of SALT 14* (pp. 110–128). Ithaca, NY.
- Graff, D. (2000). Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 20, 45–81.
- Grimshaw, J. (1991). Extended projection. Unpublished ms., Brandeis University.
- Hackl, M. (2000). Comparative quantifiers. Ph.D. thesis, Massachusetts Institute of Technology.
- Hankamer, J. (1973). Why there are two ‘than’s in English. In C. Corum, T. C. Smith-Stark, & A. Weiser (Eds.), *Proceedings of the 9th Annual Meeting of the Chicago Linguistics Society*. Chicago, IL.
- Heim, I. (1985). Notes on comparatives and related matters. Ms., University of Texas.
- Heim, I. (2000). Degree operators and scope. In B. Jackson, & T. Matthews (Eds.), *Semantics and Linguistic Theory 10* (pp. 40–64). CLC Publication Ithaca, NY.
- Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. Oxford: Blackwell.
- Hellan, L. (1981). *Towards an integrated analysis of comparatives*. Tübingen: Narr.
- Hoeksema, J. (1984). Negative polarity and the comparative. *Natural Language & Linguistic Theory*, 1, 403–434.
- Jacobson, P. (1999). Towards a variable-free semantics. *Linguistics and Philosophy*, 22, 117–184.
- Jacobson, P. (2005). Variable free semantics: The case of quantifier domain restrictions. Lecture at the Institut Jean Nicod, June 10, 2005.
- Jacobson, P. (2006). Direct compositionality and variable free semantics: Taking the surprise out of “complex variables”. Lecture at the 30th Penn Linguistics Colloquium, University of Pennsylvania, February 25, 2006.
- Kamp, H. (1975). Two theories of adjectives. In E. Keenan (Ed.), *Formal semantics of natural language* (pp. 123–155). Cambridge: Cambridge University Press.
- Kamp, H. (1981). The paradox of the heap. In U. Mönnich (Ed.), *Aspects of philosophical logic*. Dordrecht: Reidel.
- Kennedy, C. (1999). *Projecting the adjective: The syntax and semantics of gradability and comparison*. New York: Garland. (1997 UCSC Ph.D thesis).
- Kennedy, C. (2001). Polar opposition and the ontology of ‘degrees’. *Linguistics and Philosophy*, 24, 33–70.
- Kennedy, C., & McNally, L. (1999). From event structure to scale structure: Degree modification in deverbal adjectives. In T. Matthews, & D. Strolovitch (Eds.), *Semantics and linguistic theory 9* (pp. 163–180). Ithaca, NY.
- Kennedy, C., & McNally, L. (2005). Scale structure and the semantic typology of gradable predicates. *Language*, 81(2), 345–381.
- Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4, 1–45.
- Klein, E. (1991). Comparatives. In A. von Stechow, & D. Wunderlich (Eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*. (pp. 673–691) Berlin: Walter de Gruyter.
- Kyburg, A., & Morreau, M. (2000). Fitting words: Vague language in context. *Linguistics and Philosophy*, 23(6), 577–597.
- Larson, R. K. (1988). Scope and comparatives. *Linguistics and Philosophy*, 11, 1–26.
- Lasnik, P. (1999). Pragmatic halos. *Language*, 75(3), 522–551.
- Lewis, D. K. (1970). General semantics. *Synthese*, 22, 18–67.
- Lewis, D. K. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic* 8, 339–359.
- Lewis, D. K. (1983). *Philosophical papers*, Vol. I. Oxford, Great Britain: Oxford University Press.
- Ludlow, P. (1989). Implicit comparison classes. *Linguistics and Philosophy*, 12, 519–533.
- Martí, M. L. (2002). Contextual variables. Ph.D. thesis, University of Connecticut.
- McConnell-Ginet, S. (1973). Comparative constructions in English: A syntactic and semantic analysis. Ph.D. thesis, University of Rochester.
- Meier, C. (2003). The meaning of *Too*, *Enough*, and *So ... That*. *Natural Language Semantics*, 11(1), 69–107.
- Montague, R. (1974). English as a formal language. In R. Thomason (Ed.), *Formal philosophy: Selected papers* (pp. 188–221). New Haven: Yale University Press.
- Neeleman, A., Van de Koot, H., & Doetjes, J. (2004). Degree expressions. *The Linguistic Review*, 21(21), 1–66.
- Paradis, C. (2001). Adjectives and boundedness. *Cognitive Linguistics*, 12(1), 47–65.
- Pinkal, M. (1979). Semantics from different points of view. In R. Bäurle, U. Egli, & A. von Stechow (Eds.), *How to refer with vague descriptions* (pp. 32–50). Berlin: Springer-Verlag.
- Pinkal, M. (1995). *Logic and lexicon*. Dordrecht: Kluwer.

- Raffman, D. (1994). Vagueness without paradox. *The Philosophical Review*, 103, 41–74.
- Raffman, D. (1996). Vagueness and context relativity. *Philosophical Studies*, 81, 175–192.
- Rotstein, C., & Winter, Y. (2004). Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics*, 12, 259–288.
- Rusiecki, J. (1985). *On adjectives and comparison in English*. New York: Longman Linguistics Library.
- Sapir, E. (1944). Grading: A study in semantics. *Philosophy of Science*, 11, 93–116.
- Sawada, O. (2005). The possible semantic diversity of the comparative constructions in English and Japanese: A construction based approach. Paper presented at the 9th International Cognitive Linguistics Conference, Yonsei University, Korea, July 19, 2005.
- Schwarzschild, R. (2005). Measure phrases as modifiers of adjectives. *Recherches Linguistiques de Vincennes*, 35, 207–228.
- Sedivy, J., Tanenhaus, M., Chambers, C., & Carlson, G. (1999). Achieving incremental semantic interpretation through contextual representations. *Cognition*, 71, 109–147.
- Seuren, P. A. (1973). The comparative. In F. Kiefer, & N. Ruwet (Eds.), *Generative grammar in Europe* (pp. 528–564). Dordrecht: Reidel.
- Soames, S. (1999). *Understanding truth*. Oxford: Oxford University Press.
- Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23(4), 391–434.
- Stanley, J. (2002). Nominal restriction. In G. Peter, & G. Preyer (Eds.), *Logical form and language*. (pp. 365–368), Oxford: Oxford University Press.
- Stanley, J. (2003). Context, interest relativity, and the sorites. *Analysis*, (63)4, 269–280.
- Stanley, J., & Szabó, Z. (2000). On quantifier domain restriction. *Mind and Language*, 15, 219–261.
- Stechow, A. von (1984). Comparing semantic theories of comparison. *Journal of Semantics*, 3, 1–77.
- Steedman, M. (2000). *The syntactic process*. Cambridge, Mass.: MIT Press.
- Svenonius, P., & Kennedy, C. (2006). Northern Norwegian degree questions and the syntax of measurement. In M. Frascarelli (Ed.), *Phases of Interpretation*, vol. 91 of *Studies in Generative Grammar*. (pp. 133–161), The Hague: Mouton de Gruyter.
- Syrett, K., Bradley, E., Kennedy, C. & Lidz, J. (2005). Shifting standards: Children’s understanding of gradable adjectives. Paper presented at the 2005 Annual Meeting of the Linguistics Society of America, San Francisco, CA.
- Syrett, K., Lidz, J., & Kennedy, C. (2006a). Scalar structures and the semantic representations of adjectives in child language. Paper presented at the 42nd Annual Meeting of the Chicago Linguistic Society.
- Syrett, K., Lidz, J., & Kennedy, C. (2006b). The semantic typology of gradable adjectives: Experimental evidence from adult and child language. Poster presented at the 2006 Annual Meeting of the CUNY Conference on Human Sentence Processing.
- Tappenden, J. (1993). The liar and sorites paradoxes: Toward a unified treatment. *Journal of Philosophy*, 90, 551–577.
- Travis, C. (1996). Meaning’s role in truth. *Mind*, 105, 451–466.
- Unger, P. (1975). *Ignorance*. Oxford: Clarendon Press.
- Wheeler, S. (1972). Attributives and their modifiers. *Noûs*, 6(4), 310–334.
- Williamson, T. (1992). Vagueness and ignorance. *Proceedings of the Aristotelian Society*, 66, 145–162.
- Williamson, T. (1994). *Vagueness*. London: Routledge.
- Yoon, Y. (1996). Total and partial predicates and the weak and strong interpretations. *Natural Language Semantics*, 4, 217–236.