

The Relative Role of Property Type and Scale Structure in Explaining the Behavior of Gradable Adjectives

Louise McNally

Universitat Pompeu Fabra

Abstract. Kennedy (2007) proposes a semantics for positive form adjectives on which the standard for ascribing an adjective *A* makes the individuals that are *A* stand out from those that are not. To account for the differences between absolute and relative adjectives, Kennedy posits that the maximal and minimal degrees on closed scales naturally make individuals stand out in a way that degrees found away from the endpoints of a scale cannot. I argue that the ability of a degree to make individuals stand out is due less to scale structure than to the nature of the property the adjective describes. Thus, degrees that are not endpoints can behave like absolute standards as long as the application criteria for the property are clear. I relate the identifiability of such criteria to whether the property ascription can be modeled in terms of rule- vs. similarity-based classification (see e.g. Hahn and Chater, 1998).

Key words: semantics, adjectives, gradability, Sorites paradox, comparison class, classification, vagueness

1 Introduction

There is a substantial literature on the semantics of gradability that distinguishes so-called RELATIVE adjectives such as *tall* from ABSOLUTE adjectives such as *closed* (Unger, 1975; Kennedy and McNally, 2005):

- (1) a. Marta is tall.
- b. The door is closed.

Kennedy (2007) identifies three important differences between these two classes (see Rotstein and Winter, 2004, for additional relevant discussion). First, the truth of sentences such as (1a) is context-dependent while that of (1b) is not. This can be seen in the fact the addition of a modifier that makes a standard of comparison or a comparison class explicit can render the former either possibly true (as in (2a-b)) or almost certainly false (as in (2c-d)), while such an addition is not even felicitous with the latter (see (3)), indicating that a contextually-specified standard or comparison class is incompatible with the adjective.

- (2) a. Compared to her friend Andrea, Marta is tall.
- b. Marta is tall for an 11-year-old.

- c. Compared to Michael Jordan, Marta is tall.
 - d. Marta is tall for a professional basketball player.
- (3)
- a. ??Compared to Door #1, Door #2 is closed.
 - b. ??That box of cookies is closed for a box my daughter has gotten into.

Second, relative adjectives generally fail to yield crisp judgments about truth, while absolute adjectives do. As a result, the former easily give rise to the Sorites paradox, exemplified in (4). For example, if the difference in height between Marta and Andrea is very small, it is unlikely that we will judge one to be tall and the other not. More generally, assuming Premise 1 in (4a) and given Premise 2 in (4b), the general form of which is widely held to be valid for relative adjectives, the absurd conclusion in (4c) will result.

- (4)
- a. Premise 1: A 1.65-meter-tall 11-year-old is tall (for an 11-year-old).
 - b. Premise 2: If x is a tall 11-year-old and y is an 11-year-old who is 1 millimeter shorter than x , then y is a tall 11-year-old.
 - c. Conclusion: A 1.05-meter-tall 11-year old is a tall 11-year-old.

In contrast, if we open a door which is closed even the smallest amount, we can easily determine that the door will no longer be closed, and, correspondingly, if we try to recreate the paradox in (4) with *closed*, Premise 2 will fail, as Kennedy argues it fails in general for absolute adjectives.

Finally, with relative adjectives we can easily find borderline cases for which it is difficult or even impossible to decide whether the adjective truthfully holds or not. For example, if we are discussing the height of a group of children of the same age and we limit ourselves to consideration of this group, ignoring general knowledge we might have about the height of children of that age, we may have trouble judging whether a 1.50-meter-tall child is tall, if the height of the children in the group ranges from, say, 1.30 to 1.70 meters. However, no such difficulty arises with absolute adjectives like *closed*.

The contrasts between relative and absolute adjectives have two sorts of implications for a general account of the semantics of adjectives. First, as will be discussed below, the contrasting behavior of the two kinds of adjectives is unexpected on degreeless semantics of the positive form such as that proposed in Klein (1980) (though see van Rooij, To appear, for a proposed solution); in contrast, it seems less difficult to account for if we assume a semantics for adjectives that includes degrees. Second, if we assume a degree-based semantics of adjectives, these contrasts make it difficult to provide a unique, general characterization of the truth conditions for the positive form. Kennedy (2007) addresses the latter question, arguing that a unified (degree-based) semantics for the positive form is possible if we take into account the basic difference in the kinds of standard values for truthful application that are associated with each kind of adjective and properly exploit a principle of Interpretive Economy that he proposes.

The goal of this paper is to argue that, while the basic intuition behind the semantics Kennedy proposes for the positive form seems correct, the details of the analysis assign too great a role to the abstract gradability properties

of the adjectives in question. I will suggest that by focusing instead on the nature of the properties that adjectives contribute, the role that adjectives play in classifying individuals according to their manifestation of these properties, and the strategies for classification that may be involved, we can arrive at a better characterization of the relative/absolute distinction.

The structure of the paper is as follows. Section 2 presents Kennedy’s account of the absolute/relative contrasts (hereafter, the Interpretive Economy account) and some challenges to it. Section 3 relates the data discussed in Section 2 to two different classification strategies: classification by similarity, and classification by rule (see e.g. Hahn and Chater, 1998). Finally, Section 4 discusses the implications for the semantics of positive form adjectives.

2 The Interpretive Economy Account and Some Challenges

The Interpretive Economy account of relative/absolute contrasts builds on an analysis of adjectives as measure functions (see e.g. Kennedy, 1999) and on the typology of scale structure developed in Kennedy and McNally (2005). On this analysis, all gradable adjectives will denote functions from entities to degrees; specific examples appear in (5).

- (5) a. **tall**(Marta) = 1.65 meters
- b. **closed**(*lx.door*(*x*)) = 0 degrees

Kennedy and McNally argue that various linguistic phenomena are sensitive to whether the scale associated with the adjective is CLOSED, i.e. whether there are maximal or minimal values in the codomain of the measure function, or OPEN, i.e. whether there are no such values. *Tall* is an example of an open scale adjective, as in principle there is no maximal or minimal value on the height scale.¹ In contrast, *closed* is a closed scale adjective.

The denotation of an adjective is converted from a measure function into a property that can be predicated of an individual via degree morphology, which introduces a STANDARD VALUE that determines whether the property truthfully applies to its argument or not. Kennedy and McNally argue that the standard value for the truthful applicability of a gradable predicate is, like the scale itself, also subject to linguistically relevant parameterization: Specifically, it can be relative, i.e. determined contextually (typically with respect to a comparison class), or absolute, i.e. fixed at a particular value. Relative and absolute adjectives are so called because they have relative and absolute standards, respectively. The fact that an adjective like *tall* admits the expression of a comparison class, as in *tall for an 11-year-old*, is evidence that its standard is relative. In the case of absolute adjectives, Kennedy and McNally assume that there are only two possible standards: either the maximal or the minimal non-zero value on the

¹ See Kennedy and McNally (2005) for justification of the perhaps counterintuitive claim that there is no minimal value on the height scale.

scale in question. The standard is maximal just in case truthful application of the adjectival predicate entails having the property in question to a maximal degree. For example, when we try to assert that a closed door has less than the highest degree of non-aperture, we derive a contradiction, and thus can conclude that the standard for closedness is maximal (see (6a)). In contrast, a standard is minimal just in case the truthful application of the adjectival predicate only requires having the property in question to the smallest possible degree, and denying that the adjective applies, as in (6b), is incompatible with having any degree of the property in question.

- (6) a. #The door is closed, but it's slightly ajar.
 b. #Montjüic is not visible from my rooftop, but I can see a tiny part of it.

Obviously, if the standard is either maximal or minimal, adding information about a comparison class or compared individual will have no effect on interpretation, and thus *for-* and *compared to*-phrases are infelicitous with absolute adjectives.

There is a strong correlation between scale type and standard value: If absolute standards must be either maximal or minimal values on a scale, it will be impossible for an open scale adjective to have an absolute standard, since by definition such scales lack maximal and minimal values. Closed scale adjectives could in principle have either absolute or relative standards, but there seems to be a strong tendency for them to have absolute standards. Just how strong this tendency is will prove to be a crucial question.

(Kennedy, 2007, p. 17-18), building on earlier work, posits a null degree morpheme *pos* to convert the adjective into the positive form of a gradable predicate of individuals and assigns it the semantics in (7), where *g* is a variable over measure functions (of type $\langle e, d \rangle$) and '*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.'²

$$(7) \quad pos : \lambda g \lambda x. g(x) \succeq s(g)$$

But what degree allows the objects that the positive form is true of to stand out? The answer to this question must be different for relative and absolute adjectives. In the former case, Kennedy argues that the value *s* returns will depend on the context. If, in a given context, the set of individuals under consideration is such that none of them stand out with respect to any of the others in the degree to which they possess the property in question, no appropriate standard value can be chosen to differentiate among those individuals. This, Kennedy argues, is what underlies the intuitive validity of the second premise of the Sorites paradox

² The semantics in (7) differs from that in e.g. Kennedy and McNally (2005) in that it is uniform for all gradable adjectives; earlier formulations distinguished between different types of *pos* for adjectives with relative vs. absolute standards.

and the failure of relative adjectives to yield crisp judgments. Specifically, when this premise is presented and evaluated, the context is typically restricted so as to involve the comparison of just two individuals. If the difference in the degree to which they manifest the property in question is very small, the reasoning goes, that difference will not be sufficient to make one stand out with respect to the other, and thus, in the absence of any additional information, if one of the individuals is considered to have that property, the other will be as well.

In the case of absolute adjectives, we have seen that the degree that stands out and thus constitutes the standard is not determined by context. Kennedy suggests that the difference between a property holding to no degree vs. to a minimal degree, and between one holding to a non-maximal vs. a maximal degree, is sufficiently salient to make maximal and minimal degrees stand out and constitute possible standard values, even though the difference in degree is just as small as in the cases that give rise to the Sorites paradox. The difference, he suggests (Kennedy, 2007, section 4.1.), is that the transition from the absence of a property to a minimal degree of its manifestation, or from a near-maximal to maximal degree of manifestation, constitutes a ‘natural transition’ (a term he borrow from Williamson, 1992), whereas no such natural transition is obviously identifiable between degrees in the middle of a scale.

A final question that arises in this comparison of absolute and relative adjectives is why adjectives with closed scales should prefer standards that are endpoints over the sort of context-dependent standards that relative adjectives use, given that nothing about the nature of a closed scale forces the standard to be an endpoint. The answer that Kennedy proposes is that natural language follows a principle of Interpretive Economy:

- (8) **Interpretive Economy:** Maximize the contribution of the conventional meanings of the elements of a sentence to the computation of its truth conditions. (Kennedy, 2007, p. 36)

This principle is intended to guarantee that when an adjective’s scale is closed, its standard will be maximal or minimal, since the scale is presumably part of the conventional meaning of the adjective, insofar as it is derivable from the possible values of the measure function that the adjective denotes. This standard will be preferred over a standard which is determined contextually. Since relative adjectives are not associated with closed scales, there will be no such conventionally provided degree that meets the requirement of making some individuals stand out with respect to others, and there will be no choice but to choose the standard contextually.

Kennedy (ibid.) tentatively suggests that Interpretive Economy is a constraint on semantic processing and comments, ‘[t]he 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 lexi-

cons and grammars).’ Independently of whether this is demonstrably the case, Interpretive Economy can be taken to reflect the view that speakers and hearers tend to be fairly conservative in their use of language, relying on conventions that have been established in the use of language rather than innovating on a constant basis. In this latter sense the principle seems a plausible one to assume. What is less clear is the viability of the reasoning that leads from the conventionalized closed adjectival scale to the inevitable choice of an endpoint standard as an account for the contrasts underlying the relative/absolute distinction.

The crucial test cases for the Interpretive Economy account are of two kinds: 1) cases of adjectives which are (or can be) interpreted with non-endpoint standards despite having closed scales, and 2) cases of adjectives with standards that behave as absolute despite not being maximal or minimal. I begin with the first sort of case.

It is difficult to see how Interpretive Economy would allow a closed-scale adjective to have a standard which is not an endpoint, unless it is simply a default principle. But such adjectives do exist: a good example is *familiar*. In (9) we find two instances of this adjective. It is very difficult to see how these instances could be interpreted differently from each other, particularly since the second sentence seems intended to defeat the scalar implicature generated by the first that the familiarity property is not held to a higher degree. The only difference is that in one case the standard is established with respect to a comparison class (with *very* offering evidence for this; see Kennedy and McNally, 2005), while in the other, it is identified with respect to the number of things the student has to be familiar with. In this latter case, the standard is a minimal degree of familiarity; *completely* indicates that the scale associated with the adjective is closed (Hay et al., 1999).

- (9) For a student who has just moved here, she is very familiar with the class routines and her teachers’ expectations. In fact, she’s completely familiar with them.

On the Interpretive Economy account, the adjective contributes the same measure function in the two cases: a function from entities to degrees of familiarity whose maximal and minimal values are determined by the volume of stuff that corresponds to the *with* argument. The unanswered question is why this conventional input can be ignored and a comparison class used to establish the standard on some occasions. Some kind of contextually licensed override of the economy principle would have to be possible, or perhaps one could argue that the meaning of *familiar* has been conventionalized in such a way that the standard is no longer a maximal value. Either way, we do not gain much insight into how and why such deviations are possible.

The Interpretive Economy account, as developed, also does not predict the existence of adjectives whose standards behave as if they were absolute without being endpoints on a scale. This is not due to the principle itself, but rather to the ancillary claim that maximal and minimal degrees are fundamentally different from degrees which are not at the endpoints of scales in providing natural

transitions. I now turn to a couple of examples that challenge this ancillary claim. These counterexamples do not call into question either Interpretive Economy as a principle, or the semantics for the positive form in (7), but they do shift the burden of explanation from scale structure to the specific nature of the adjectival properties themselves.

The first of these examples involves the adjective *full*. Kennedy and McNally (2005) argue that *full* is an absolute adjective with a maximal standard, and Syrett (2007) presents psycholinguistic evidence in support of that claim, based among other things on the fact that when presented with the two pictures in Figure 1, adult subjects consistently considered a request to indicate ‘the full [jar]’ infelicitous.



Fig. 1. Stimulus for the test item ‘The full one’ from Syrett (2007), Appendix E.

However, the facts are more subtle than these results indicate at first. Foppolo and Panzeri (2010) present experimental data that indicates that subjects’ intuitions about what counts as the standard for the equivalent of *full* and certain other absolute adjectives in Italian is sensitive to the type of object being ascribed the property. Indeed, Kennedy and McNally (2005) already observed that speakers are in some cases willing to accept that a container is full even it is not full to the maximal degree. For example, when one is served a full glass of soda or beer, the beverage rarely reaches the top. Kennedy and McNally nonetheless maintained the claim that the standard is the maximal degree and offered two possible explanations for speakers’ intuitions. One is that is that when a nearly-full glass is claimed to be full, the predication is, strictly speaking, false, but interlocutors are willing to speak loosely if the predication is sufficiently close to true for the purposes of the context (see Lasersohn, 1999). The other is that the granularity of measurement might be made coarse enough in some cases so as to allow an almost full glass to count as full. However, neither of these explanations is plausible when we consider cases such a glass of wine. Normally a wine glass is considered full if it is filled to about half of its capacity with wine. It is extremely

difficult to imagine that in such cases we are taking the glass's maximal capacity as the standard but either speaking loosely or applying a very coarse granularity of measurement.

The defender of the maximal standard could argue as an alternative that we simply ignore part of the volume of the glass, thus conserving the maximality of the standard in the form of a degree that amounts to reaching the fill line for the glass, which might be below the glass's maximal capacity. This could be formalized by treating *full* as something like a function from container types to measure functions whose codomains have possibly different specific maximal values. Thus, a sentence like (10) could be accounted for by positing different choices of container type (e.g. the most generic form of glass in the first case; a glass in which wine is served in the second) for each of the uses of the adjective.

- (10) The wine glass is half full; therefore, it is completely full as far as this restaurant is concerned.

This would preserve the Interpretive Economy account, but at the price of conceding that the measure function that *full* denotes depends on the specific sort of object to which it is applied as well as factors such as the function the container is fulfilling. In other words, the standard would depend indirectly on the context. Moreover, this analysis carries additional commitments. For instance, (11) should be true when the glass is filled to the fill line:

- (11) This wine glass is completely full; it cannot be fuller.

But in fact it is difficult to deny that a functionally full wine glass can be made fuller, even under the interpretation that takes into account a 'fill-line' standard. It seems more promising to rethink the conditions on the use of the degree modifier *completely*, which might pose the biggest challenge to the claim that the standard is not a maximal value. Instead of requiring that *x is completely A* iff *A(x)* is the maximal value on the volume scale, as proposed in Hay et al. (1999), we could simply require it to be true just in case *A(x)* is that degree which corresponds to the volume that would be occupied after a properly completed filling event involving the container in question. Such a degree might be unique for each kind of container, but it is far from obvious that it is maximal.³

Another challenge for the claim that the standard for *full* need not be a maximal value comes from the fact that, even when applied to the wine glass, it behaves like the standard of an absolute adjective. First, it is not compatible with *for-* or *compared to*-phrases:

- (12) a. ??Compared to the glass on the table, this glass is full.
b. ??This glass is full for a wine glass.

³ In fact, in some cases it might even exceed the volume capacity of the container by some amount that might be difficult to quantify precisely, as Ede Zimmermann (p.c.) observes might be the case with sake glasses.

Second it licenses crisp judgments and, correspondingly, reasoning about *full* does not fall into the pattern of the second premise of the Sorites paradox. Finally, it does not yield borderline cases within the limits of our ability to measure. If the standard is not a maximal degree, we are then left with the question as to what differentiates this non-maximal standard for *full* and the non-maximal standards for typical relative adjectives. I will return to this question in the next section.

An even clearer example of a standard that challenges the claim that all absolute standards are endpoints is that of color terms on a color-extension reading, where the gradable property makes reference to the amount of the object in question that has the color property. Consider, for example, the shirt in Figure 2:



Fig. 2. Boy's Gray Camo T-shirt (from <http://www.teamcamogear.com>)

Both (13a) and (13b) are arguably true statements about this shirt.

- (13) a. This shirt is gray, but not completely gray.
- b. This shirt is not black.

This suggests that in order for a color extension property to hold, the degree to which it holds must be more than minimal (otherwise, (13b) would be false) but less than maximal (otherwise (13a) would be false). Moreover, in contrast to the case of *full*, it also does not seem plausible to argue that we ignore those parts of the shirt that are not gray and that the standard is maximal with respect to the parts that are gray: (14) sounds like a contradiction.

- (14) The shirt is completely gray, but it's not completely gray.

Nonetheless, as with *full*, this standard does not behave like a relative standard: The adjective does not allow overt mention of a comparison class (see (15a)), nor

does it admit degree modification by *very*, which Kennedy and McNally (2005) argue is compatible with relative adjectives in general ((15b)).⁴

- (15) a. ??This shirt is gray for a patterned t-shirt.
b. ??This shirt is very gray.

Rather, the standard behaves like an absolute standard, with the difference that, instead of having to be maximal or minimal, it seems to be fixed as to a degree such that the color gray predominates.

Though this standard behaves like an absolute standard in not being sensitive to a comparison class, we should ask whether it gives rise to crisp judgments and fails to yield borderline cases as well. In the case of crisp judgments, the answer is, I think, positive in principle. While it may be difficult to put a fixed value on the amount of color that counts as predominating (presumably it is above 50% but what exact degree it is might vary depending on e.g. where the color is distributed on the object in question), it is possible to imagine that a very small reduction in degree of color extension might make the difference between a color predominating or not, and thus the adjective applying truthfully or not. The fact that this might happen means that the second premise of the Sorites' paradox cannot hold as a rule for these adjectives. Borderline cases should also not arise, insofar as it should be possible for speakers to decide for any given case whether a color predominates in an object or does not. If it does not predominate, the color term will not apply truthfully.

Thus, the assumption that only endpoints on a scale can serve as natural transitions and, independently of reference to a comparison class, make some individuals stand out with respect to others, seems too strong. Some other explanation for the properties of absolute standards must be identified in order for the semantics in (7) to be maintainable.

Let us now turn very briefly to relative adjectives. Kennedy's claim is that the function *s* in (7) chooses a standard that makes the right individuals stand out as having the property in question. He eventually identifies standing out with being on the upper end of a natural transition. Since, on his view, no degree on an open scale provides a natural transition, a comparison class must be appealed to in order to identify the standard. However, nothing is said about how the standard is actually selected. Consider again *tall*. In what sense does the shortest member of a group of tall individuals stand out against the tallest member of the group of *non-tall* individuals, or lie on the upper end of a natural transition from non-tall to tall? This is the question that the Sorites paradox confronts us with. I want to suggest that we should look for the answer by attacking the problem from a different direction.

One fact that sometimes gets lost in discussions of the Sorites paradox is that one of the functions of adjectives is to group individuals according to the way they manifest a given quality (e.g. in the case of height, we have not only groups for *tall* and *short*, but also for some uses of e.g. *tiny* or *gigantic*). We

⁴ Of course, (ib) could be used to describe the intensity of the color of a shirt. The '??' indicates anomaly as a description of color extension.

might therefore consider the possibility that it is not the function s that ‘chooses a standard...in such a way as to ensure that the objects that the positive form is true of “stand out”,’ but rather the possible groupings of objects manifesting a given quality in different degrees that determine the standards for the adjectives that make reference to that quality.⁵ In the following section, I suggest that by approaching the determination of the standard in this way, we can see the sense in which the standards eventually chosen for both absolute and relative adjectives can be considered natural transitions, even though the naturalness of those standards is not always evident by looking at the scale itself.

3 Classification and Standards

Let us take as our starting point the claim that the truthful application of an adjectival predicate requires the denotation of the adjective’s argument be sortable into the category of objects for which the predicate holds, taking into account the range of relevant predicates that might characterize that individual along the quality or dimension of interest. We can then ask whether there are relevant differences in the way this sorting or classification task might proceed for relative vs. absolute adjectives, and whether these differences correlate with the contrasts presented at the beginning of this paper. The answer, I suggest, is positive: specifically, the truthful predication of relative adjectives can be insightfully modeled as CLASSIFICATION BY SIMILARITY, whereas the truthful predication of absolute adjectives can be modeled as CLASSIFICATION BY RULE.

Hahn and Chater (1998) propose two criteria for distinguishing similarity- vs. rule-based reasoning, including in particular categorization or classification.⁶ First, they claim that rule-based classification depends on a strict matching between the classification criterion/a and the relevant properties of the object being classified. In contrast, similarity-based classification requires only a partial match. Second, they maintain that rule-based classification involves comparing a representation associated with a specific individual (for instance, one concerning the degree of fullness of a specific glass) against a representation that is more abstract (e.g. a degree of fullness for glasses in general), whereas similarity-based classification involves comparing a representation of a specific individual or property of that individual against another representation of an equally specific individual or one or more of its properties.

Consider now the prototypical absolute adjectives. When the standard for such adjectives is a maximal or minimal degree, it is trivial to see how the decision about whether they apply to their arguments could be formulated as a simple rule which would not require comparison with any specific individuals.

⁵ Barker’s (2002) analysis of adjectives such as *stupid* might be considered to represent the spirit of this perspective on the way the standard is determined.

⁶ Space limitations preclude a full discussion of the difficulties involved in distinguishing these two kinds of classification; the reader is referred to Hahn and Chater’s article for details and a defense of the position that they *are* meaningfully distinguishable.

To know whether a (generic) container is full, we need only know how much of its volume is occupied. If all of it is, the container is full; if not, it is not. To know whether a door is open, we need only check whether it has any aperture at all, and so on.

The view that the property contributed by an absolute adjective is ascribed via rule directly accounts for the behavior of these adjectives described above. First, since no comparison needs to be made to specific individuals, no comparison class is called for; thus, the examples in (3), repeated in (16), should be odd.⁷

- (16) a. ??Compared to Door #1, Door #2 is closed.
b. ??That box of cookies is closed for a box my daughter has gotten into.

Second, if absolute adjectival properties are ascribed via rule, we can account for their failure to yield the Sorites paradox. Classification by rule, as described, requires an exact match of the classification condition associated with the property. If this condition is not precisely met, there is no reason to think that the adjectival predicate will truthfully apply. Of course, this does not exclude the possibility of a certain variability in the granularity of measurement or other criteria of precision that will be applied in order to decide whether there is a match and that the individual counts as having the property; however, the imprecision in these cases will turn out to be more circumscribed than in the case of classification on the basis of similarity.⁸ Finally, in the case of absolute adjectives, indeterminacy or borderline cases will only arise to the extent that the application criteria for the rule are not fully defined or, in the case of a measurable property, the granularity of measurement is coarse.

Crucially, there is no reason to think that rule-based classification will be possible only for cases where the standard for a gradable adjective is either maximal or minimal. We only need a precise and principled way to identify the standard degree in question in order to be able to formulate the rule we need. For example, there is nothing problematic about a rule to the effect that a wine goblet is full if half of its volume is occupied, or that a beer glass is full if it contains 33 centiliters of liquid. Such rules can be applied without any consideration of a specific comparison class - they only require reference to the type of container involved insofar as that plays a role in the rule for determining fullness. Such criteria not only permit but require an exact match in the manifestation of the property in the individual in question. Thus, even though the standard is neither maximal nor minimal, we expect these adjectives to behave like other absolute adjectives. Similarly, a rule for determining color extension can be formulated, as suggested above, in terms of perceptual predominance of the color.

⁷ Note that direct comparison to a specific individual or class of individuals is not the same as reference to a specific *sort* of individual to clarify which variant of a rule for property ascription will apply.

⁸ See van Rooij (To appear) for a fuller discussion of the relation between granularity and vagueness.

While classification by rule is thus clearly possible in ascribing properties denoted by absolute adjectives, we should ask whether anything in principle forces it. The answer is clearly negative. The use of *full* to characterize density of volume occupation, as in (17a-b), is a good example.

- (17) a. For a Friday, the dentist’s schedule is very full.
 b. Compared to the last box you packed, this one is very full.
 c. The dentist’s schedule/The box is completely full.

Clearly there is a maximal value for the volume occupation of a schedule or a box (see (17c)), but in (17a-b) the standard is determined by a comparison class or compared individual. Since most relative uses of *full* that I have identified involve characterizing density of volume occupation, we can perhaps look for an answer to the question of what factors lead to the maintenance of an absolute standard or the introduction of a relative one by asking specifically whether there is something about a property like density of volume occupation that lends itself to ascription via similarity-based reasoning.⁹ Here, I think perceptual factors come into play: one thing we might expect from our unmodified, non-technical vocabulary is that it be usable without the help of measuring tools.

Density of volume occupation is at least partially independent of whether the contents of a container come into contact with the physical limits of that container. It thus may be difficult to measure precisely without some kind of measurement tool when the physical limits of a container are reached by the contents, though we can be sure that a container whose limits are not reached will not be fully occupied. Containers whose physical limits *are* reached by contents with a certain perceptible density will share important properties in common with a container whose volume is completely occupied, and they will be decidedly distinct from any container whose physical limits are not in contact with the contents. These similarities and differences may be sufficiently useful and salient for speakers to classify the former containers as full and the latter as not full, even though in such a situation it may nonetheless be clear that the similar containers are not identical. In such cases, the addition of information about comparison class or specific compared individuals may be added to improve the precision of the property ascription.

Now consider prototypical relative adjectives, again using *tall* as an example. Why should classification by rule with such adjectives be impossible? Here Kennedy’s intuition about the relevance of the lack of endpoints on the scale associated with the adjective seems exactly on target. Rule-based classification involving non-maximal/minimal standards can work for adjectives interpreted with respect to a closed scale because such adjectives describe properties that can be held to proportional degrees. Since at least certain proportions are perceptually salient and can be easily estimated without knowing absolute values, it can be comparatively easy to know when a given property is held to a specific

⁹ Relative uses of *full* where density is arguably not at issue, like *a very full glass of water*, can also be found; I would give an explanation for them that is similar in spirit, if different in detail.

proportion. Nothing like this will be possible when applying an adjective that contributes a gradable property characterized by a single, unbounded dimension, uncorrelated with any other easily and consistently observable characteristic. In the case of height, the only plausible rule we could apply would be to stipulate a specific height value as the standard. But such a value can only be expressed in an arbitrarily chosen measurement unit. Ascribing tallness to someone under such circumstances would depend on our ability to carry out this kind of measurement, and in the context of everyday language use this simply might not be feasible often enough.¹⁰

In contrast, if we have previously identified exemplars of tall and short individuals,¹¹ classifying any given third individual as tall or short on the basis of relative similarity to these exemplars is entirely feasible. In particular, it does not necessarily entail measuring in any precise way the heights of any of the individuals in question; to identify an individual as tall, it is sufficient to be able to judge that individual as more similar to the tall exemplar(s) than to the short one(s).¹² Nor should we worry about where the necessary exemplars to do this classification would come from: note, for example, that an entire sub-genre of children's literature is devoted to conveying certain qualities in terms of opposites, obviously facilitating familiarity with reliable pairs of exemplars on the basis of which children can learn to ascribe the relevant properties to new individuals.

If relative adjectives are ascribed on the basis of similarity rather than rule, we can immediately explain their behavior. First, the ascription of the property contributed by the adjective will depend crucially on a comparison class or compared individual because the classification involves ascribing the property on the basis of a comparison to representations of specific individuals; the job of the comparison class is to provide these individuals, most crucially, the exemplars that will serve as the basis for the classification. Second, we can explain the intuitive validity of Premise 2 of the Sorites paradox: This premise conflicts with a basic principle of similarity-based classification, which is that one classifies by maximizing within-class similarity and between-class distance. Finally, borderline cases can arise because some individuals may prove to be equally similar to the exemplars of the classes under consideration and thus difficult or impossible to classify in a non-arbitrary way. We can also now characterize the sense in

¹⁰ Of course, an explicit standard can be introduced via a measure phrase or a comparative expression (e.g. *taller than Max*). Unsurprisingly, comparatives do not allow modification by comparison clauses, do not give rise to the Sorites paradox, and do not have borderline cases.

¹¹ For simplicity, I will make this point assuming that with height there is just a binary classification into tall and short, but of course additional classes are arguably called for, such as neither-tall-nor-short, midget-sized, etc. These can be easily incorporated into the analysis.

¹² This is just the basic sort of algorithm that is used in the simplest forms of clustering, a standard similarity-based classification technique. See e.g. Gärdenfors (2000) for application of notions from clustering-based approaches to classification to a theory of conceptualization, specifically his theory of Conceptual Spaces.

which the standard for relative adjectives is a natural transition: it will be that degree which marks the boundary between the classes that result from grouping the individuals in question according to their similarity.

The reader might also be concerned that, on this view, we would end up with two very different kinds of satisfaction conditions for gradable adjectives, which might roughly be characterized as follows, setting aside degrees for the moment and assuming that we take adjectives to denote properties, and assuming that multiple properties (such as tallness or shortness) can be related to a given quality or dimension (such as height):

- (18) $\mathbf{AdjP}_{absolute}(x)$ is true in context C iff x manifests the property contributed by \mathbf{AdjP} as required in C .
- (19) $\mathbf{AdjP}_{relative}(x)$ is true in context C and relative to a comparison class K iff x is more similar to the exemplar from K for the property contributed by \mathbf{AdjP} than it is to the exemplar for any other property under consideration for classifying x with respect to the quality or dimension in question.

However, note that there is no need to assign the two sorts of adjectives to different logical types, and thus the fact that the satisfaction conditions are different should not be a cause for concern.

Summarizing, the position defended here is that abstract scale structure properties are only indirectly related to the contrasts in the behavior of absolute vs. relative adjectives. What is crucial is the possibility of establishing clear applicability conditions for the rule. Various factors can facilitate or impede the establishment of such conditions. I have suggested that one important factor is the ease with which the degree that constitutes the standard can be perceived; maximal and minimal standards are a special case of such degrees, but not the only one.

I now turn briefly to the implications of the preceding discussion for the semantics of the positive form.

4 Implications for the Positive Form

Kennedy (2007) takes the contrasting behavior of relative and absolute adjectives as an argument against a degreeless semantics for adjectives on which they denote simple properties of individuals, and specifically against the analysis defended in Klein (1980). Klein's semantics treats adjectives as functions which assign individuals to positive extensions, negative extensions, and extension gaps, where extensions are defined relative to a domain D – effectively equivalent to a comparison class – of cardinality greater than or equal to 2. For any well-defined comparison class, Klein assumes that the positive and negative extensions of the adjective with respect to that class must not be empty; he defines the semantics in this way so that quantification over comparison classes can be used to induce the ordering needed to support a semantics for comparatives. Specifically, on

this view *x is Adj-er than y* iff there is a comparison class for which *x* falls into the positive extension of the adjective and *y* does not. But as Kennedy notes:

[it is not clear] 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. (Kennedy, 2007, p. 41)

In other words, since comparison classes are irrelevant for, and indeed infelicitous with, the positive form of absolute adjectives, it should not be possible for an object to be in the positive extension of a relative adjective with respect to one comparison class and in the negative extension with respect to another, and if this is not possible, then it will not be possible to correctly apply Klein's analysis of comparatives to absolute adjectives. However, this criticism might be overcome if there was some other way to induce the ordering needed to support comparatives involving absolute adjectives. Bale (2006) and van Rooij (To appear) propose such alternatives.

The starting point of the discussion in the previous section was the degree-based semantics in (7), repeated in (20), so we should therefore reconsider it in light of that discussion:

$$(20) \quad pos : \lambda g \lambda x. g(x) \succeq s(g)$$

This sort of degree-based semantics have been criticized on the grounds that it is very abstract, that there is never or almost never any morphological manifestation of *pos*, and that it would appear that the semantics of the positive is effectively defined in terms of a comparison relation (see e.g. Klein, 1980; Bale, 2006; van Rooij, To appear). But let us set aside these concerns and focus on the key issue, which is whether the standard-fixing function *s* can be defined in a unified way for both absolute and relative adjectives.

We have, in a sense, found a way to characterize what it means for a standard to be a natural transition both for absolute and for relative adjectives. In the former case, it is that degree which the rule for truthful application of the adjective effectively makes reference to. In the latter, it is that degree which marks the boundary between the groupings that are derived from a similarity-based classification of a set of individuals according to the conventionalized labels we have for identifying the different manifestations of the quality or dimension in question. Nonetheless, at a deeper level, it is difficult to see how the value returned by *s* could be characterized in any truly unified terms across absolute and relative adjectives other than as 'the degree that makes the adjectival predicate truthfully hold.' Moreover, in the case of relative adjectives, I do not see any way to derive this standard degree except in an *a posteriori* fashion on the basis of the way the members of the comparison class are sorted in any given context.

But if this is the case, it would seem that, as mentioned above, the identification of the standard presupposes that we are able to successfully use the adjective. This is of course not a problem for a strictly formal account of adjective semantics, but it lends support to the criticisms of such a semantics as a useful model of our semantic competence.

5 Conclusion

The discussion of gradable properties in the formal semantics literature has been heavily conditioned by a focus on orderings, scales, and standards, rather than on the nature of the properties themselves or the general role that adjectives play in categorization. While we must not lose sight of the need to derive orderings or precise measurements to support a semantics for comparative and related constructions, I have argued here on the basis of some new and overlooked data that a focus on the nature of the properties described by adjectives, as well as the classification strategies that might successfully model their ascription, can be particularly useful in providing a better understanding of the relative/absolute distinction.

Acknowledgments

I am very grateful to Christopher Kennedy, not only for many years of collaboration on the semantics of adjectives but also very specifically for giving me the opportunity to first develop the ideas presented in this paper, in the form of a commentary on his ‘Vagueness and Grammar’ paper at the Chicago Workshop on Scalar Meaning held in May 2006. I would also like to thank Carla Umbach, who first pointed out to me the work by Hahn and Chater, to Kristen Syrett both for comments and for permission to reproduce the material in (9), and to the editors of this volume for very helpful criticism. Finally, thanks to the organizers of the 2007 Amsterdam Colloquium and the 2008 DGfS Workshop on Comparison and Similarity for giving me the chance to present this work there, to the audiences there and at CIMEC in Rovereto for their feedback, and to the organizers of the 2009 ESSLLI Vagueness and Communication workshop, who perhaps expected to hear a version of this paper at the workshop but didn’t. The work reported here was supported by the Spanish Ministry of Science and Innovation under grant HUM2007-60599/FILO and by a Fundació ICREA Acadèmia award.

Bibliography

- Bale, A. (2006). *The Universal Scale and the Semantics of Comparison*. Ph.D. thesis, McGill University.
- Barker, C. (2002). ‘The dynamics of vagueness’. *Linguistics and Philosophy*, 25:1–36.
- Foppolo, F. and Panzeri, F. (2010). ‘Do children know when their room counts as clean?’ Ms., University of Milano-Bicocca.
- Gärdenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*. Cambridge, MA: MIT Press.
- Hahn, U. and Chater, N. (1998). ‘Similarity and rules: distinct? exhaustive? empirically distinguishable?’ *Cognition*, 65:197–230.
- Hay, J., Kennedy, C., and Levin, B. (1999). ‘Scale structure underlies telicity in degree achievements’. In T. Matthews and D. Strolovitch (eds.), *Proceedings of SALT IX*. Ithaca, NY: CLC Publications, 127–144.
- Kennedy, C. (1999). *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. New York: Garland. (1997 UCSC Ph.D thesis).
- (2007). ‘Vagueness and grammar: The semantics of relative and absolute gradable adjectives’. *Linguistics and Philosophy*, 30:1–45.
- Kennedy, C. and McNally, L. (2005). ‘Scale structure and the semantic typology of gradable predicates’. *Language*, 81:345–381.
- Klein, E. (1980). ‘A semantics for positive and comparative adjectives’. *Linguistics and Philosophy*, 4:1–45.
- Lasnik, P. (1999). ‘Pragmatic halos’. *Language*, 75:522–551.
- Rotstein, C. and Winter, Y. (2004). ‘Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers’. *Natural Language Semantics*, 12:259–288.
- Syrett, K. (2007). *Learning about the Structure of Scales: Adverbial Modification and the Acquisition of the Semantics of Gradable Adjectives*. Ph.D. thesis, Northwestern University.
- Unger, P. (1975). *Ignorance*. Oxford: Clarendon Press.
- van Rooij, R. (To appear). ‘Vagueness and linguistics’. In G. Ronzitti (ed.), *The Vagueness Handbook*. Berlin/Heidelberg: Springer.
- Williamson, T. (1992). ‘Vagueness and ignorance’. *Proceedings of the Aristotelian Society*, 66:145–162.