Chris Kennedy
University of Chicago

Seminar on Vagueness
23 May, 2006

Vagueness, Imprecision and Scale Structure

1 Context and the positive form

Last week we discussed the dynamic analysis of the positive form of vague gradable adjectives developed by Barker (2002). Simplifying a bit, the analysis assigns the meaning in (1) to *pos*, where **d** is a function that maps a context and a measure function to a degree.

(1) $\lambda g \lambda x \lambda c.g(x) \succeq \mathbf{d}(c)(g)$

This analysis leads us to expect two kinds of contextual variability in the standard of comparison: variability based on properties of c (e.g., the domain of discourse) and variability based on properties of g (e.g., the domain of the measure function). Both kinds of variability seem attested.

Moreover, this analysis ends up predicting the sort of 'dual use' of vague expressions that Barker discusses: a sentence like (2a) can be construed either as 'about Betty' (the desriptive use) or as 'about the context' (the metalinguistic use).

(2) a. Betty is tall.
b. $\lambda c.\mathbf{tall}(\mathbf{betty}) \succeq \mathbf{d}(c)(\mathbf{tall})$

This analysis is also going to be able to support a number of different explanations of the Sorites Paradox and borderline cases/boundarylessness.

(3) a. Steve Nash is short (for a basketball player).
b. Any basketball player who is 1 mm taller than a short basketball player is also short (for a bb player).
c. Shaquille O'Neal is short (for a bb player).

(4) a. $\mathbf{short}(\mathbf{nash}) \succeq \mathbf{d}(c_0)(\mathbf{short})$
b. $\forall x, y, c[[[\mathbf{short}(y) \succ \mathbf{d}(c)(\mathbf{short})] \wedge [\mathbf{tall}(x) - \mathbf{tall}(\mathbf{y}) = \mathbf{1mm}]] \rightarrow$
$[\mathbf{short}(x) \succ \mathbf{d}(c)(\mathbf{short})]]$
c. $\mathbf{short}(\mathbf{shaq}) \succeq \mathbf{d}(c_0)(\mathbf{short})$

On pretty much all the analyses we discussed, (4b) turns out to be false; the interesting part is why we think it's true. And if we're wright about the semantics of *pos*, then I think our options remain open.

Well, maybe: now that we have an explicit semantics, let's finally decide on the significance of 'crisp judgments'. Consider a context in which Betty is just a little bit taller than Abe. All of the examples in (5) are bad, while those in (6) are ok.

(5) a. #Betty is tall compared to Abe.
b. #Betty is the tall one.
c. A: What counts as tall around here? Is Abe tall?
B: No, but Betty is.

(6) a. Betty is taller than Abe.
b. Betty is the taller one.

Let's say that the function of *compared to* and the definite article *the* are similar in these examples: they ensure that the context relative to which the standard of comparison is computed is one that consists only of Abe and Betty. For example, *compared to* might do something like the following:

(7) $[\![$x compared to y$]\!] = \lambda f_{\langle s,t \rangle} \lambda c. f(c[x,y])$, where $c[x,y]$ is a context just like c except that the domain of discourse consists of just x and y.

It seems to me that something like Graff's analysis does the best job here. A supervaluation analysis fails miserably; it is not clear that Williamson's margin of error principe should apply if we're using two arbitrary objects to fix the standard; and given that Raffman wants to allow for categorical judgments, it seems reasonable that if we're invited to consider a context in which there are only two objects *that are known to differ in their respective amounts of the relevant feature*, then we ought to put the boundary between them.

If, on the other hand, **d** is the 'significant degree' function (for a context c and a measure function g), then we're in business: our interests in efficiency, etc. should be just as active in this context as in any other, so the cost of discriminating between two objects that have very similar degrees of some property should outweigh the benefits of doing so, and we end up treating them the same — this *despite* the pressure to impose a partitioning on the domain! (Optimality is everywhere!)

So let's take (8) to be the denotation of *pos* (for the rest of the day, at least!), where **s** is the significant degree function.

(8) $\lambda g \lambda x \lambda c. g(x) \succeq \mathbf{s}(c)(g)$

Note also that this kind of contextual analysis should solve both of the problems that Stanley (2003) raised: because the context variable is bound, we can invoke a 'sloppy identity' analysis of ellipsis to respond to his first worry, and we can index it to the context of utterance to handle is second worry.

2 Vagueness and grammar

2.1 Absolute gradable adjectives

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

(9)
a. The gold is impure.
b. The table is wet.
c. The door is open.
d. The rod is bent.

Maximum standard absolute adjectives such as those in (10) require their arguments to posses a maximal degree of the property in question.

(10)
a. The platinum is pure.
b. The floor is dry.
c. The door is closed.
d. The rod is straight.

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

(11)
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.

(12) 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:

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

THE PROBLEM: if absolute adjectives have the same semantic type as relative adjectives, their interpretations in the positive form are unexpected.

(14) a. $[\![pos]\!]([\![\text{impure}]\!]) = \lambda x \lambda c.\mathbf{impure}(x) \succeq \mathbf{s}(c)(\mathbf{impure})$
b. $[\![pos]\!]([\![\text{pure}]\!]) = \lambda x \lambda c.\mathbf{pure}(x) \succeq \mathbf{s}(c)(\mathbf{pure})$

(14a) is true of an object if it has a contextually significant degree of impurity, and (14b) is true of an object if it has a contextually significant degree of purity. These are not accurate characterizations of the truth conditions of these predicates.

Or is it? (15) can be felicitously used to describe a situation in which only a very few people show up to a film in a very large movie theater.

(15) The theater is empty tonight.

MY RESPONSE: This is IMPRECISION, which is distinct from vagueness. In particular, the former is about use, while the latter is about meaning. We'll get into this in detail in the second part of class today, but first I want to go through a few quick arguments to establish the relative/absolute distinction. (We have seen these before.)

2.1.1 Natural precisifications

Imprecision is most clearly illustrated by measure phrases. Both (16a) and (16b) could be used to describe a rod whose actual length falls somewhere close to 10 meters, in a range that is itself be subject to contextual variation.

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

However, (16a) but not (16b) allows for what Pinkal (1995, pp. 99-100) calls NATURAL PRECISIFICATIONS: we can construct a context in which *10 meters long* distinguishes between objects based on very slight differences in length (e.g., a scientific experiment or a construction project); it is difficult (if not impossible) to do the same for *long*.

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

Absolute adjectives allow natural precisifications. If I am a detective in search of a violent criminal, and I think he is in the movie theater described above, I would judge a projectionist who responds to my question *Is anyone in the theater tonight?* with (15) to be lying.

2.1.2 Entailments

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

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

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

(18a) is true only if the floor has some degree of wetness: if it had zero wetness, then it could not be wetter than the countertop. This satisfies **??**, entailing (18b).

Similarly, in order for (19a) to be true, it must be the case that the countertop is not maximally dry. It therefore follows from **??** that the countertop is not dry.

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

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

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

This is expected: the fact that one object exceeds another with respect to some relative property tells us nothing about how the objects stand in relation to a contextually significant amount of the relevant property.

A potential problem?

(22) Both A and B are dry, though A is even drier than B.

2.1.3 The Sorites Paradox

In (23), for example, the second premise is quite naturally judged to be false, bypassing the paradox.

(23) 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.

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

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

P2 is again easily judged false: one way for a quantity of water to have fewer contaminants than impure water would be to have no contaminants at all, but in that case the water would no longer be impure!

2.2 The polysemy of the positive form

The descriptive conclusion to draw from these facts seems to be that the positive form is polysemous:

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

a. $\lambda g \lambda x \lambda c. g(x) \succeq \mathbf{s}(c)(g)$ (when g is a relative A)

b. $\lambda g \lambda x. g(x) \succ min(\text{SCALE}(g))$ (when g is an absolute minimum A)

c. $\lambda g \lambda x. g(x) = max(\text{SCALE}(g))$ (when g is an absolute maxumum A)

This raises the following questions:

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

I think the answer to the first question is to some extent a functional one: the meaning of the positive form is a mapping from a gradable adjective (a measure function) to a property of individuals that relates the degree to which an individual possesses the property measured by the adjective to a reference point that is computed as a function of the meaning of the adjective (the standard of comparison).

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

- The maximal and minimal degrees on a scale are fixed values that can be computed strictly on the basis of the function expressed by the adjective (its range).
- A contextually significant degree of the property expressed by the adjective supports a meaningful partitioning of the domain of the adjective.
- Others? Maybe: *prototypes* (color terms); *natural transitions* (*early/late, flat/sharp, fast/slow* (of clocks), etc.); *tropes* (evaluative adjectives, etc.).

We still have to deal with the second question. For example, why does *long* combine only with (25a), deriving a relative interpretation, while *bent* combines only with (25b) and *straight* only with (25c), deriving minimum and maximum absolute interpretations, respectively? Why don't absolute adjectives like *straight* and *bent* also combine with (25a), deriving relative interpretations? More generally:

- What factors are responsible for the fact that only relative adjectives are vague?

The crucial generalizations (Kennedy and McNally 2005):

(26) a. Gradable adjectives that use open scales have relative interpretations in the postive form.

b. Gradable adjectives that use closed scales have absolute interpretations in the positive form.

(26a) follows directly from the interaction of scale structure and the polysemous semantics of *pos*: open scales lack maximal and minimal values, so open scale adjectives are incompatible with maximum and minimum standards, leaving only the relative interpretation.

(26b), however, does not follow. A closed scale structure may restrict the range of interpretations that *pos* can have, but it does not uniquely determine a particular interpretation. In particular, there is no semantic incompability between (totally or partially) closed scales and a relative interpretation of *pos*.

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

2.3 Interpretive economy

Features of the conventional meaning of a gradable adjective can have consequences for the truth conditions of the positive form. One example of this is the following principle discussed in Kennedy to appear.

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

The generalization about the relation between closed scales and absolute standards is similar in form:

(28) *'Range Dependence'* If a gradable adjective expresses a function to a closed scale, then the standard of comparison must correspond to a maximal or minimal element of the scale.

I would like to suggest that the similarity between these two descriptive generalizations is not accidental, but rather that both follow from a more general principle of INTERPRETIVE ECONOMY:

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

The basic idea is that truth conditions should be computed on the basis of the conventional meanings of the expressions of a sentence (or logical form) to the extent possible, allowing for use of contextual information in the calculation of truth conditions only as a last resort, when conventional meaning is insufficient.

2.3.1 Restricted Domains

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

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

Since *large* is an open scale adjective, only the relative meaning of *pos* is an option for the interpretations of the predicates:

(31) a. $\lambda x \lambda c.\mathbf{large}(x) \succeq \mathbf{s}(c)(\mathbf{large})$
b. $\lambda x : (x) \lambda c.\mathbf{large}_{mammal}(x) \succeq \mathbf{s}(c)(\mathbf{large}_{mammal})$

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

- Significance is an inherently relative notion: in order to determine the minimal degree that represents a significant amount of size, you need a property on which to base this judgment.

- The conventional meaning of the adjective in (31a) is a function from objects that have sizes to sizes. The mere property of 'having size' is insufficient to determine what counts as a significant degree of size, so some other property must be recovered from the context. The property expressed by the modified nominal is highly salient and so is a likely candidate, though as we have seen this is only a preference, not a requirement.
- The conventional meaning of the (modified) adjective in (31b) is a function from mammals to sizes. The property of being a mammal provides a basis on which to calculate a significant degree of size; since this property is part of the conventional meaning of the adjective, *Interpretive Economy* dictates that it must be used in computing the truth conditions of the positive form.

Note that context sensitivity is not entirely eliminated even in this case, since what counts as significant is also a function of the interests and expectations of the participants in the discourse (Graff 2000). But use of the restricted domain in (31b) maximizes the contribution of the conventional meaning of the constituents of the predicate to the calculation of its truth conditions, in accord with *Interpretive Economy*. That is, it is enough to look just at g; we don't also need to look at c.

2.3.2 Partially Closed scales

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

Consider the adjectives *bent* and *straight*, which have lower-closed and upper-closed scales, respectively. For *bent*, both of the interpretations in (32) are in principle possible.

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

The truth conditions of (32b) are computed strictly on the basis of the conventional meaning of *bent* (its range), while (32a) introduces context dependence via the **s** function. *Interpretive Economy* therefore selects the absolute interpretation in (32b).

For *straight*, the choice is between the relative interpretation in (33a) and the maximum standard intepretation in (33b).

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

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

2.3.3 Totally Closed Scales

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

The antonyms *opaque* and *transparent* verify this prediction. (34a) can be felicitously utterered at the point at which I have almost reached 100% of tint, demonstrating both that *opaque* can have a maximum standard (I am denying that the glass is completely opaque) and that *transparent* can have a minimum standard (partial transparency).

(34) CONTEXT: Manipulating a device that changes the degree of tint of a car window from 0% (completely transparent) to 100% (completely opaque)
 a. The glass is almost opaque, but not quite. It's still transparent.
 b. The glass is almost transparent, but not quite. It's still opaque.

Likewise, (34b) can be used to describe the reverse situation: one in which I have dialed down almost to 0% of tint. Here *transparent* has a maximum standard (complete transparency) and *opaque* has a minimum standard (partial opacity).

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

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

(35) CONTEXT: On the starship *Enterprise*
 a. If the airlock is open, the cabin will depressurize.
 b. The ship can't be taken out of the station until the space door is open.

Anyone who does not understand (35a) to be a warning that partial opening of the airlock will result in depressurization, is a danger to the ship and crew, and so is a helmsman who fails to understand (35b) as an prohibition against trying to leave the station before the space door is completely open.

Similarly, (36a) is most naturally understood to describe a line of troops in which some soldiers are exposed to enemy fire (as pointed out to me by Mark Richard).

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

However, the reason that ABC got in trouble with the FCC over (36b) (during the halftime show of Superbowl XXXVIII) was because Janet Jackson's breast was exposed all the way to the nipple, not because there was some skin visible.

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

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

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

In fact, the choice of standard is systematic. Focusing on the case of deverbal adjectives derived from accomplishment verbs, Kennedy and McNally (2005) observe that the standard associated with a total scale deverbal adjective is a function of the role played by the argument of the adjective in the event described by the verb:

- If the argument corresponds to the incremental theme of the source verb, then the adjectival form has a maximum standard.
- If not, then the adjectival form has a minimum standard.

For illustration, consider *loaded.* (38a-b) show that a goal argument gets a minimum standard, but an incremental theme argument gets a maximum standard.

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

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

- In the case of an incremental theme argument, the relevant event is not completed unless the argument has been totally affected by the verb; this will only be the case if it has a maximum degree of the property expressed by the adjective. A minimum standard interpretation is therefore filtered out.
- With a non-incremental theme argument, completion of the event is consistent with a situation in which the argument merely has a non-zero degree of the property expressed by the adjective. Either a minimum or maximum standard interpretation should be possible.

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

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

Here what is being denied is that the truck is fully loaded, not that it has some degree of loadedness.

Most adjectives with totally closed scales appear to be deverbal, though not all are related to accomplishment verbs. The adjectives in (40), for example, are related to achievements, but the core of the explanation outlined above carries over directly.

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

The verbal source of the adjectives in (40a) (*open, uncover, expose*) names the initiation point of an event, while the verbal source of those in (40b) (*close, cover, hide*) names the culmination point.

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

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

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

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

3 Imprecision

Lasersohn (1999) discusses three kinds of 'loose talk', and ways of making such cases of imprecision precise.

(42) a. Mary arrived at 3 o'clock.
b. The townspeople are awake.
c. The line is straight (linear).

According to Lasersohn, all of the examples in (42) are strictly speaking false in many natural contexts of use; in such contexts, we are just using them imprecisely. Evidence comes from contradictions like those in (43).

(43) a. ??Although the townspeople are awake, some of them are asleep.
b. ??Although Line A is straight, it is not as straight as Line B.

Compare:

(44) a. Although pretty much all the townspeople are awake, some of them are asleep.
b. Although Line A is long, it is not as long as Line B.

Precision can be forced by SLACK REGULATORS like *exactly, all* and *perfectly*, as in (45).

(45) a. Mary arrived at exactly 3 o'clock.
b. All the townspeople are awake.
c. The line is perfectly straight (linear).

Lasersohn (1999, p. 535): 'We should distinguish, I think, between authentic semantic vagueness, in which the extension of a predicate really does not have well-defined borders, and mere pragmatic looseness of speech, in which we allow speakers a certain leeway even in the use of predicates whose extensions are quite sharply defined.'

The core idea: We can use sentences like (42) informatively even when they are false as long as they are 'close enough to the truth' so as not to obscure pragmatically relevant distinctions.

We determine whether a proposition is 'close enough to true' but looking at its PRAGMATIC HALO.

- For any expression α, the pragmatic halo of α is a set of objects of the same semantic type as $[\![\alpha]\!]$ that differ from it in pragmatically ignorable ways.
- The pragmatic halo of α is under an ordering which ensures that $[\![\alpha]\!]$ is at its center. (The farther away you go, the less similar to α you become.)
- Halos are built up compositionally.
- The size of a halo (the degree of imprecision) is a matter of context.
- Slack regulators restrict halos to their centers, but also introduce halos themselves. (Almost every term has a halo; see below.)

A proposition is 'close enough to true' if its pragmatic halo is non-empty (contains at least one possible world). So even if Unger is right about nothing ever being flat, we can now understand why *The road is flat* is informative.

References

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

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

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

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

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

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

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

Stanley, Jason. 2003. Context, interest relativity, and the sorites. *Analysis* 63.4:269–280.