Chris Kennedy
University of Chicago

Seminar on Vagueness
18 April, 2006

Partial Models, Supervaluation and Degrees of Truth

1 Review: Gradable predicates and vagueness

Last week we considered two hypotheses about gradable adjective meaning. On *the Degree Analysis*, GPs map objects to degrees on a scale (they are type $\langle e, d \rangle$ or $\langle d, et \rangle$), and are 'converted' to properties of individuals by degree morphology.

(1) a. [DegP Deg AP] b. [AP DegP A]

(2) $[\![Deg]\!]([\![AP]\!]) = \lambda x.D(\mathbf{g}(x))$, where $\mathbf{g}$ is a function from objects to degrees on the scale expressed by $[\![\text{AP}]\!]$ and D is a property of degrees determined by $[\![\text{Deg}]\!]$.

(3) Possible values of D:
   a. *measure two meters*
   b. *exceed Kim's height*
   c. *be less than the maximal degree to which Kim was happy yesterday*
   d. ...

As these examples show, comparison in this system involves relations between degrees, and results in a predicate that is not vague. The basic architecture allows for vagueness via appropriate specification of D. In particular, if D is something like (4), then vagueness could arise from the standard computation function, whatever that is.

(4) *exceed a standard of comparison computed on the basis of the domain of g*

Alternatively, the standard computation *could* in principle be precise (the mean, average, whatever), and vagueness could come about in some other way. More on this later. What is important to keep in mind is that on this analysis:

- Gradability and vagueness are formally distinct.
- Gradable adjectives are not vague.
- Vagueness is a function (directly or indirectly) of the semantics of the positive form (degree morphology or type-shifting rule).

At the most general level, what a semantics of degree gives us is a 'gradability parameter': a way of directly encoding in the compositional semantics the fact that particular expressions support orderings on their domains.

The second hypothesis treats vague predicates just like other property denoting expressions (type $\langle e, t \rangle$), with one exception: they denote *context dependent, partial* functions from objects to truth values, rather than (non-context dependent) total functions. In any context of use, it is necessary to fix the positive and negative extensions of a vague predicate (the things it is definitely true and false of), and it is possible (in fact normal) that some things in its domain will fall in an *extension gap* — these are the things it is neither definitely true nor definitely false of.

Exploring what exactly this means from a semantic and logical perspective will be the focus of today's class. But before we do that, it is important to recognize some important distinctions between this analysis and the degree analysis:

- Gradability and vagueness are directly connected, because comparison involves consideration of alternative ways of putting objects in the positive/negative extension of a vague predicate.
- Gradable adjecives are inherently vague.
- Vagueness is a function of the basic type of a gradable predicate (partiality) and the definition of the model for interpreting such predicates.

Before we get into the details, I want to throw out a few questions to keep in mind along the way:

1. On the degree analysis, should anything that shows the properties of vagueness (borderline cases, sorites susceptibility) have a degree semantics, or is there more than one way to be vague? If the latter, what are the options, and how can we distinguish 'degree vagueness' from other kinds of vagueness?
2. On the partial function analysis, could any category whose members are of type $\langle e, t\rangle$ potentially include vague members, i.e., members that are partial in the relevant way? Does every vague term support comparison?

In particular, I want us to think about the predictions/expectations of the two analyses for instances of non-adjectival vague predicates.

2 Supervaluations: The basic idea

Let's go back to one of the fundamental properties of vague predicates: the existence of borderline cases. (5a) is clearly true; (5b) is clearly false; (5c) is indeterminate. (I know, because my son asked me whether this was true, and I wasn't sure how to answer, other than by saying 'It depends...'!)

(5) a. $[\![\text{Mercury is close to the sun}]\!] = 1$
b. $[\![\text{Pluto is close to the sun}]\!] = 0$
c. $[\![\text{The earth is close to the sun}]\!] = ??$

However, despite the uncertainty about (5c), we seem to have pretty clear intuitions about (6a-b): the former is a logical truth, while the latter is a contradiction.

(6) a. The earth is close to the sun or far from the sun.
b. The earth is close to the sun and far from the sun.

Assume that $[\![\textit{the earth is far from the sun}]\!] = \neg[\![\textit{the earth is close to the sun}]\!]$:

(7) a. $p \vee \neg p$ (Law of the Excluded Middle)
b. $p \wedge \neg p$

Fine 1975: Logical relations can hold between sentences without clear truth values in a particular context of utterance. According to Fine, this tells us that truth valuations for propositions built out of logical connectives should be based not on particular contexts, but on an appropriate space of interpretations in which all 'gaps' in truth values have been filled — the ways that vague expressions can be *precisified*. Here are the crucial components of the theory:

1. *Specification space:* A partially ordered set of points corresponding to different ways of specifying the predicates in the language; at each point, every proposition is assigned true, false or nothing according to an 'intuitive' valuation (*appropriate specification*; cf. Klein's *Consistency Postlate*).

2. *Completeability:* Any point can be extended to a point at which every proposition is assigned a truth value.
   (a) FIDELITY: Truth values at complete points are 1 or 0
   (b) STABILITY: Definite truth values are preserved under extension
3. *Supertruth:* A proposition is true (or false) at a partial specification iff it is true (false) at all complete extensions.

In short, the reason that we have clear intuitions about (6a-b) is is because for any ways of making things more precise, we're always going to end up with a situation where (7a) holds for any p and (7b) fails to hold, regardless of whether p has a truth value at the beginning. (7a) is *supertrue* (and (7b) is superfalse).

This theory gets us out of the Sorites Paradox, too: (8b) is (super) false because on any complete and admissible specification there is going to be a sharp boundary between the heavy and non-heavy stones.

(8) a. A 100 kilo stone is heavy.
b. Any stone that weighs 1 gram less than a heavy one is heavy.
c. #A 1 gram stone is heavy.

According to Fine, we tend to accept (8b) for two reasons. First, the value of the cutoff point appears to be arbitrary. But is this right? Or is it that the very existence of a cutoff point seems arbitrary? I'm not exactly sure....

Perhaps the answer is related to the second point: in different complete specifications, the cutoff point will be in different places. In principle, it could be an infinite number of places, if we allow an infinite domain (though it will always have to respect 'admissibility'). There need be no determinate fact about where it is; moreover, on the view outlined here, nothing about this is *linguistic*, though presumably linguistic factors (e.g., the meanings of the expressions used) could come into play, and maybe even play a role in determining admissible completions of the specification space.

What is the meaning of a vague predicate on this view? The following seems to be the core idea (p. 277):

(9) "Let the *actual* meaning of a simple predicate, say, be what helps determine its instances and counterinstances. Let its *potential* meaning consist of the possibilities for making it more precise. Then the point is that the meaning of an expression is a product of both its actual and potential meaning. In understanding a language one has thereby understood how it can be made more precise; one has understood ... the possibilities for its growth."

These remarks are clearly applicable to non-gradable predicates. Are they applicable to gradable ones? Perhaps to some of them:

(10) a. *empty/full, straight, dry, pure, ...*
b. *bent, wet, impure, ...*
c. *red, blue, green, ...*
d. *happy, sad, intelligent, ...*

They are applicable to something like *have a height/weight/speed/cost that exceeds the norm/average for objects of the relevant type*, but are they applicable to *tall/short, heavy/light, fast/slow, expensive/inexpensive* and so forth???

Perhaps this type of framework should be imposed *on top* of a basic degree semantics? There's nothing inconsistent about that, though it does seem a bit redundant.... Perhaps we need to take a closer look at the actual analysis of gradable predicates in this sort of system.

3 Comparatives and degrees of truth

Kamp 1975: Comparatives show that we want to analyze gradable adjectives as expressions of type $\langle e,t\rangle$, not $\langle et,et\rangle$. On the latter view, (11a) is underlyingly (11b), where the value of NP (a function of type $\langle e,t\rangle$) is recovered from context.

(11) a. Kim is tall.
 b. Kim is a tall NP.

This means that the truth conditions of (12a) are (12b).

(12) a. Kim is taller than Lee.
 b. $\exists f \in D_{et}[tall(f)(kim) \wedge \neg tall(f)(lee)]$

But what if Kim and Lee differ in height to a very small degree? Kamp argues that there is no coherent value for f in this situation, yet the comparative is perfectly fine: it supports CRISP JUDGMENTS. So we need another option, and according to him, the logical structure that we need in order to deal with partial models, specifications and supervaluations provides the basis for a semantics of comparatives that does the right thing, and moreover provides us with a useful notion of 'degree of truth' — without actually introducing degrees into our ontology. Here are the basic components:

1. Assume (as with Fine) partial models in which some predicates denote complete functions to {0,1} and others denote partial functions with positive extensions, negative extensions and extension gaps. This allows for vagueness in particular contexts of utterance in the way that we saw above.
2. Relate a partial model to all classical models that are:
   (a) completions of the original one, or
   (b) based on variants of the original one with different positive/negative extensions and extension gaps;
   (c) in all cases, the models satisfy (what Klein calls) the Consistency Postulate: if a is Aer than b in a base model, then there is no completion/variant in which b is in the positive extension of A but a is not.
3. The *degree of truth* of a proposition p in a model M is the conditional probability of certain completions of M (those based on the 'ground model') over the set of all completions. (See his discussion of *a is heavy* on p. 139.)

The basic account of borderline cases and the Sorites Paradox is the same as in Fine's account. What Kamp adds is a way of linking vagueness to gradability/comparison (or probably more accurately, vice-versa). Specifically, he give us two options for the semantics of comparatives to choose from:

(13) $[\![x \textit{ is more A than } t]\!]^M = 1$ iff
 a. $\{M' \mid [\![x \textit{ is } A]\!]^{M'} = 1 \wedge M' \textit{ is a completion of } M\} \supset$
 $\{M' \mid [\![y \textit{ is } A]\!]^{M'} = 1 \wedge M' \textit{ is a completion of } M\}$
 b. $d([\![x \textit{ is } A]\!]^M) \succ d([\![y \textit{ is } A]\!]^M)$, where d is the 'degree of truth' function

Either view gets around the 'crisp judgment' problem, which is a very good result!

Kamp argues that 'multidimensional' adjectives like *skillful, large* etc. argue for the former over the latter approach. Specifically, he's worried about how to evaluate (14) in the various situations

(14) Smith is cleverer than Jones.
  a. Smith is much better at solving mathematical problems than Jones.
  b. Jones is more quick-witted than Smith.

Kamp's worry: the disparity in (14a) ought to make this sentence true, because *Smith is clever* is going to have a greater degree of truth than *Jones is clever*.

Maybe easier to judge:

(15) Chicago is larger than Los Angeles.
  a. Chicago has a greater population than Los Angeles.
  b. Los Angeles covers more square miles than Chicago.

How do we handle these within a degree-based analysis? **Measure function polysemy**

Probabilities do get used for the analysis of degree modifiers, though here we seem to be very close to a true degree semantics.

What determines the 'core meaning' of a gradable predicate in a particular context of utterance? For Kamp, this includes both the positive/negative extensions and the dimension, and here he just bites the bullet and says that the context of utterance does this. He also says that *linguistic* context plays a central role, letting him explain the role of modified nominals in providing 'comparison classes' (though he is careful to point out that they do not have to do this, as we saw last week).

But again, what is the 'core' meaning of a vague predicate, and to what extent do we expect them to show uniform behavior?

Two issues to consider:

- It follows from Kamp's system that anything that is vague should be gradable. What does this lead us to expect about nouns like *table*, and so forth? (Kamp's discussion of this is on pp. 148-9.)

- It is less clear whether it follows from Kamp's system that anything that is gradable is vague.

The facts, however, <u>are</u> clear: not all gradable expressions are vague. ABSOLUTE GRADABLE ADJECTIVES provide the crucial evidence for this, which means going back to last week's handout....