

To appear in *Linguistics and Philosophy*. Version of 30 May, 2001, final version with proof corrections included.

The dynamics of vagueness¹

Chris Barker <barker@ucsd.edu>
Linguistics, University of California, San Diego

1. Introduction

Many people have studied how vague predicates depend on context for their interpretation, but few have studied in detail how a use of a vague predicate affects the context against which other expressions get evaluated. In recent years, however, various dynamic theories of context update have made considerable progress in describing and explaining complex phenomena such as presupposition and anaphora. In this paper I will argue that taking an explicitly dynamic perspective on vagueness can lead to new insights into the nature of vagueness in general and the semantics of gradable adjectives in particular.

In pursuit of this goal, I will provide an explicit formal analysis of how a use of a vague gradable adjective (e.g., *tall*) affects shared knowledge in a developing discourse. My starting point is the claim that asserting and accepting a token of the sentence *Feynman is tall* can resolve some portion of the mutual uncertainty associated with the applicability of the predicate *tall*. More specifically, it eliminates from further consideration the possibility that the vague standard of absolute tallness might be greater than the maximal degree of Feynman's height. I develop this idea within a dynamic update semantics in the tradition of Stalnaker, Karttunen, Heim, and many others.

1.1. Descriptive versus metalinguistic modes of use

One advantage of adopting a dynamic perspective is a clearer picture of the situations in which adjectives have a descriptive versus a metalinguistic use. Normally, (1) will be used in order to add to the common ground new information concerning Feynman's height:

(1) Feynman is tall.

But (1) has another mode of use. Imagine that we are at a party. Perhaps Feynman stands before us a short distance away, drinking punch and thinking about dancing; in any case, the exact degree to which Feynman is tall is common knowledge. You ask me what counts as tall in my country. "Well," I say, "around here, ..." and I continue by uttering (1). This is not a descriptive use in the usual sense. I have not provided any new information about the world, or at least no new information about Feynman's height. In fact, assuming that *tall* means roughly 'having a maximal degree of

¹Svetlana Godjevac, Beth Levin, and Chris Kennedy provided invaluable advice over many drafts. In addition to the anonymous *L&P* referees and audiences at UCSD and SALT 10, thanks also to David Beaver, Daniel Büring, Dan Brassil, David Dowty, Donka Farkas, Jean Mark Gawron, Delia Graff, Jorge Hankamer, Larry Horn, Alice Kyburg, William A. Ladusaw, Peter N. Lasnik, Sally McConnell-Ginet, Louise McNally, Paul Portner, Jason Stanley, Gina Taranto, and Timothy Williamson.

height greater than a certain contextually-supplied standard', I haven't even provided you with any new information about the truth conditions of the word *tall*. All I have done is given you guidance concerning what the prevailing relevant standard for tallness happens to be in our community; in particular, that standard must be no greater than Feynman's maximal degree of height. The context update effect of accepting (1) would be to eliminate from further consideration some candidates for the standard of tallness. My purpose in uttering (1) under such circumstances would be nothing more than to communicate something about how to use a certain word appropriately—it would be a metalinguistic use. In fact, I will argue that metalinguistic update is by no means pathological or exceptional, but part of the normal update potential of most vague predicates and present to one degree or another in most (but not all!) uses.

1.2. Two contrasting case studies: comparatives and infinitive complements

Sections 3 and 4 extend the basic analysis to two detailed case studies investigating comparatives and certain infinitival-taking adjectives. The discussion of comparatives treats measure phrases (*six feet tall*), comparatives (*Bill is taller than Feynman*), and degree modifiers (*very tall*), with special attention to higher-order vagueness (*definitely tall*). Measure phrases and comparatives have no metalinguistic side-effects: declaring that Bill is six feet tall, or that Bill is six inches taller than Feynman, reveals nothing about what counts as tall. On the analysis here, the impossibility of using measure phrases or comparatives to negotiate vague standards follows from the truth conditions of the constructions.

The second case study involves adjectives that take infinitival complements (*happy to see you*, *proud to be your friend*, *lucky to survive*). In general, adjectives behave quite differently semantically when they take an infinitival complement.

- (2) a. Feynman is stupid. ABSOLUTE
b. Feynman is stupid to dance like that. RELATIVE

In (2a), Feynman is habitually stupid, or disposed or likely to behave stupidly. In (2b), Feynman's stupidity is limited to his participation in a specific dancing event. Certainly neither sentence entails the other: Feynman might very well be stupid to dance wildly, in which case (2b) is true, at the same time he is a Nobel laureate, in which case (2a) may very well be false.

Getting the update potential of relative uses of adjectives right can explain a striking phenomenon:

- (3) a. Feynman wanted to be proud/stupid.
b. Feynman wanted to be [proud to dance like that].
c. *Feynman wanted to be [stupid to dance like that].

As Wilkinson (1970:432) observes and as (3) illustrates, when some adjectives such as *stupid*, *lucky*, *rude*, *wise*, *fortunate*, etc., take an infinitival complement, they can no longer be embedded beneath a control predicate like *want* (except on an irrelevant purpose-clause interpretation). My explanation for this unembeddability effect depends on showing that in these infinitival uses, what

is special about adjectives like *stupid* is that they have no entailments whatsoever (beyond what they presuppose) except for their update effect on vague standards. That is, I will argue that relative uses of stupid-type adjectives have only a metalinguistic mode of use.

One larger goal of this study is to lend support to the idea that complex semantic behavior often follows automatically from getting local truth conditions right. An especially clear example of this general strategy is Heim's (1983; 1992) claim that intricate details of presupposition projection follow automatically in a dynamic framework from the truth conditions of the component expressions. My specific goal, then, is to show how providing a dynamic account of the truth conditions of gradable adjectives leads to an explanation of the update behavior of comparatives and degree modifiers on the one hand, and of unembeddability on the other.

2. The dynamics of gradable adjectives: basic analysis

2.1. Background

The dynamic approach developed here is compatible with Bosch's (1983:190) idea that vagueness is a case of 'incomplete definition', i.e., incomplete acquisition of the precise meaning of a predicate. A similar idea is elaborated at length in Williamson's (1994) proposal that vagueness is (p. 3) 'an epistemic phenomenon... In cases of unclarity, statements remain true or false, but speakers of the language have no way of knowing which.' Thus for Williamson, a use of a vague predicate reduces ignorance. Williamson (pages 205–215), however, rejects making a systematic distinction between knowledge concerning the content of a discourse versus knowledge concerning the discourse itself. I will insist on being more specific: part of the ignorance associated with a use of a vague predicate is uncertainty about the applicability of a word. That is, it is ignorance pertaining to the state of the discourse itself—in other words, it is purely linguistic ignorance.

Therefore I will distinguish three types of entailment: **presuppositions** are what must already be accepted as true in order for a use of an expression that triggers those presuppositions to be appropriate; **sharpening** is what happens to the status of vague predicates as discourse constrains possible ways of resolving that vagueness²; and (normal) **descriptive** entailments, of course, are at-issue entailments other than sharpening entailments and presuppositions that follow directly from the truth conditions of the utterance. Thus a typical assertion of the sentence *John's bald uncle died* presupposes that John has an uncle, sharpens our notion of what counts as bald, and has among its descriptive entailments that John's uncle at some point ceased to be alive.

The observation that a use of a vague predicate affects subsequent context is not new by any means. Kamp (1975:149), Klein (1980:14), and Pinkal (1989:223) all mention in passing that asserting a sentence containing a vague predicate causes the context to change in certain ways, but none spell out in detail the mechanism of those changes. At more length, Parikh (1994) examines vagueness as a problem for effective communication. He describes something very much like what I call

²'Sharpening' as used here corresponds to Pinkal's (e.g., (1983)) non-dynamic notion 'precization'. See also Williamson's (1999) notion of a sharpened language.

sharpening, and he advocates a view of meaning that resembles the dynamic analysis developed here. It is difficult to tell for sure, however, since Parikh's discussion, though cogent, is relatively informal. Eikmeyer and Rieser (1983) and Ballweg (1983:70) are much more explicit about the way in which the interpretation of a vague expression depends on context and how the context must change as discourse progresses. More recently, Kyburg and Morreau (1996; 2000) offer an accommodation-based formal model for updating the common ground to reflect the appropriate degree of sharpening of a vague predicate following an utterance. Kyburg and Morreau's theory describes one aspect of "negotiation about how to use vague words", and thus in its broad strokes is very much in the same spirit as this paper.³ However, Kyburg and Morreau explicitly note that they are not proposing an algorithm for computing updated values for vague predicates, but rather are only placing constraints on what such an algorithm "ought to compute". What is still missing, and what this paper attempts to provide, is an explicit account of precisely how the context change potential of a use of a vague expression follows directly from its compositional contribution to truth conditions.

2.2. *Formal preliminaries*

The formal approach I will take to modeling context update and presupposition projection is in the tradition of Stalnaker (1973 et seq.), Karttunen (1974), Heim (1982, 1983), Groenendijk and Stokhof (1990, 1991), and many others. The basic idea of these dynamic approaches to meaning is that the mutual assumptions of the interlocutors can be modeled as a set of alternatives, where each alternative corresponds to one way the world might be.

Most dynamic theories that discuss anaphora provide a separate formal mechanism for tracking variable assignments. For instance, in Heim (1983), a context consists of a set of ordered pairs $\langle g, w \rangle$ where g is an assignment function mapping variables to entities and w is a possible world. In recognition of the fact that such contexts seem to contain more than just worlds, they are sometimes called 'information states'. In fact, Kamp (1988) argues that sets of possible worlds are not rich enough to model anaphoric relations appropriately, so that a separate mechanism is not only convenient, but required. But as Stalnaker (1998) explains, if the worlds in a context are all worlds in which the discourse itself is taking place, then each world contains within it a complete version of the discourse itself—including the complete set of anaphoric relations that are part of (that world's version of) the discourse.

Thus the need for a separate bookkeeping mechanism for variable assignments can be reconciled with a pure possible-worlds conception of update by providing a function \mathbf{g} that maps a world to the assignment function for the (counterpart of the) discourse in which the sentence under evaluation occurs. Such a function will be defined only for worlds in which the discourse in question is taking place, but it is quite natural (in fact, unavoidable) to assume that speakers will presuppose that they are participants in a discourse.

Analogously, \mathbf{d} will be a function that maps a world onto the delineation that characterizes the

³I became aware of this work only during final revisions, and the two theories were developed independently of one another. There are significant differences in framework and analysis: Kyburg and Morreau explicitly allow a limited degree of belief revision, whereas my model provides only for monotonic increase in information; in addition, their treatment of higher-order vagueness is quite different from the one below in section 3.4.

vague predicates in use in the discourse of that world. A DELINEATION, in turn (Lewis (1970)), is a function from adjective meanings to degrees.

- | | |
|---------------------------------------|--|
| (4) <small>FUNCTION ON WORLDS</small> | <small>SEMANTIC TYPE OF THE VALUE</small> |
| g | assignment function (maps variables to individuals) |
| d | delineation (maps gradable adjective meanings to a degree) |

Thus given a context C , each candidate world $c \in C$ specifies how things might be along three dimensions simultaneously: who might be salient, how vague predicates might be made precise, and all the other facts that determine what might be the case. For instance, if c is a candidate world in a discourse context, then $\mathbf{g}(c)$ is the variable assignment function associated with c , so that $\mathbf{g}(c)(x)$ yields the referent of the variable x in c . Similarly, $\mathbf{d}(c)$ is the delineation function associated with c , and $\mathbf{d}(c)(\llbracket tall \rrbracket)$ yields the standard of absolute tallness in c .

As usual in dynamic systems, propositional expressions denote context update functions, i.e., a function that takes a context as an argument and returns an updated context for a value. If ϕ is an expression in English and C and C' are contexts, $\llbracket \phi \rrbracket(\langle C, C' \rangle)$ holds just in case C' happens to be a legitimate way of updating C with the information in ϕ . I will say that $\llbracket \phi \rrbracket$ is the UPDATE FUNCTION denoted by ϕ , and given that $\llbracket \phi \rrbracket$ is a function, it will sometimes be convenient to write $\llbracket \phi \rrbracket(C) = C'$. Thanks in part to the fact that contexts are sets of worlds (rather than information states), all of the update functions considered below are filters, so that if $\llbracket \phi \rrbracket(\langle C, C' \rangle)$ holds then $C' \subseteq C$. Therefore if $\llbracket \phi \rrbracket$ is an update function I will sometimes slightly abuse notation and write $C \subseteq \llbracket \phi \rrbracket$ instead of $\llbracket \phi \rrbracket(C) = C$, and similarly for $c \in C$ write $c \in \llbracket \phi \rrbracket$ instead of $\llbracket \phi \rrbracket(\{c\}) = \{c\}$. Thus if **rain** is the set of worlds in which it is raining, $\llbracket It is raining \rrbracket = \lambda C. \{c \in C : c \in \mathbf{rain}\}$.

2.3. Descriptive and metalinguistic modes

Being more specific about how a use of a vague adjective updates context requires talking about degrees. Following many modern treatments of gradable adjectives (see von Stechow (1984); Klein (1991); Kennedy (1997, in press) for surveys), I will assume that gradable adjectives in all of their uses express relations involving individuals and degrees. Therefore I will express the denotation of *tall* in terms of the logical constant relation **tall**, where $\mathbf{tall}(d, \mathbf{b})$ will be the set of worlds in which Bill is tall at least to degree d .

But what is the ontological status of a degree? Many theories assume that degrees are (isomorphic to) real numbers (see discussion in Hamann (1991), Klein (1991), or Kennedy (in press)), which is natural enough for adjectives such as *tall*, but less obviously appropriate for adjectives whose degrees are more abstract, such as *stupid*. For present purposes, it will suffice to assume that the set of degrees D is partially ordered by the relation \leq , and that gradable adjectives are monotonic with respect to this ordering (see, e.g., Klein 1991:684). That is, if an individual is tall (or stupid) to degree d , then that individual is tall (or stupid) to any degree d' where $d' \leq d$. This assumption is what guarantees, among other things, that if Bill is tall and Feynman is not tall, then Bill must be taller than Feynman, and so on.⁴

⁴Two technical considerations: Kennedy (in press) argues that degrees must be intervals, not points, in order to account for a

Absolute uses of adjectives depend on a contextually-supplied standard, which is where the delineation function introduced in (4) comes into play. An absolute use such as *Feynman is tall* will be true just in case there is a degree to which Feynman is tall that exceeds the prevailing standard for absolute tallness. In general, we have:

$$(5) \quad \llbracket tall \rrbracket = \lambda x \lambda C. \{c \in C : c \in \mathbf{tall}(\mathbf{d}(c)(\llbracket tall \rrbracket), x)\}$$

This dynamic denotation takes an individual as an argument and returns an update function. The denotation in (5) says that an individual x will count as absolutely tall in a world c just in case x is tall (at least) to the degree identified by the delineation function associated with c .⁵

In order to be able to put *tall* in a sentence, assume that the relevant use of *be* is semantically transparent, i.e., $\llbracket be \rrbracket = \lambda X.X$, and that the [NP VP] construction gets the normal generalized quantifier construal, i.e., $\llbracket [NP VP] \rrbracket = \llbracket NP \rrbracket(\llbracket VP \rrbracket)$. As usual, a name will denote a principal ultrafilter, so that $\llbracket Feynman \rrbracket = \lambda P.P(\mathbf{f})$, where \mathbf{f} is the individual named Feynman, and we have

- (6) a. Feynman is tall.
 b. $\llbracket Feynman \text{ is tall} \rrbracket = \lambda C. \{c \in C : c \in \mathbf{tall}(\mathbf{d}(c)(\llbracket tall \rrbracket), \mathbf{f})\}$

This update function will exclude worlds based on the standard of tallness in that world and Feynman's height in that world.

In a purely descriptive use, the prevailing standard of tallness is common knowledge. More technically, this means that the range of tallness standards that \mathbf{d} assigns to the candidates in the common ground are sufficiently close to one another not to affect the truth of the assertions under discussion. To see what this might be like, consider an utterance of (6) in the presence of a context containing the following candidates:

phenomenon he calls cross-polar anomaly. The semantics given here is neutral with respect to this issue; however, if degrees are points, then care must be taken to distinguish between upward monotonicity for positive adjectives (*tall*, *wise*, etc.) versus downward monotonicity for their antonyms (*short* and *stupid*, respectively), which would fall out automatically on Kennedy's assumptions. Second, Sapir (1944) points out that if we conceive of colors as regions on a continuous spectrum, a gradable adjective like *green* must be bounded on both sides: a color is green only if it sufficiently un-blue and and also sufficiently un-yellow. Therefore Lewis (1970) suggests associating each adjective with a tuple of standards, each element of which marks the cutoff point along a different dimension (see also McConnell-Ginet (1973)). We could sidestep both issues by allowing the delineation function to return a set of degrees, the set of all degrees that qualify as tall or the set of all degrees that qualify as green. Such sets may be intervals, as claimed by Kennedy. However, making do with one degree per adjective as a standard for resolving vagueness will suffice for our purposes here.

⁵The equation in (5) mentions $\llbracket tall \rrbracket$ on both sides of the equal sign, which might seem to make the definition circular in a pernicious way. However, since delineations map adjective meanings onto degrees (rather than, say, onto some type of function), it is easy to find a suitable value for $\llbracket tall \rrbracket$ that satisfies the constraint expressed by (5).

(7)	c	$\mathbf{d}(c)(\llbracket tall \rrbracket)$	$\iota(\max(\lambda d.\mathbf{tall}(d, \mathbf{f})))$	$c \in \mathbf{tall}(\mathbf{d}(c)(\llbracket tall \rrbracket), \mathbf{f})$
	c_1	180	183	yes
	c_2	180	181	yes
	c_3	180	179	no
	c_4	180	177	no

Let c_1, \dots, c_4 represent four candidate possible worlds that differ relevantly only in the maximal degree to which Feynman is tall. Let's assume that Feynman is tallest in c_1 (in which he is 183 centimeters tall), progressively shorter in c_2 and c_3 , and shortest in c_4 . The values in the column labelled $\mathbf{d}(c)(\llbracket tall \rrbracket)$ are identical, which means that variations in the standard of tallness in the four worlds will play no role in the update outcome. Only in worlds c_1 and c_2 is Feynman tall enough to qualify as tall in the absolute sense, since only in c_1 and c_2 is the maximal degree to which Feynman is tall greater than 180 centimeters. This means that only these two candidates survive update with the denotation of the sentence *Feynman is tall*. Thus if c_1, \dots, c_4 is the context at the point in a discourse immediately before the sentence in question is processed, c_1, c_2 will be the updated context: $\llbracket Feynman is tall \rrbracket(\{c_1, c_2, c_3, c_4\}) = \{c_1, c_2\}$. This scenario constitutes a purely descriptive use of (6), and the sole update effect is to add information to the context about how tall Feynman must be without sharpening the notion of what counts as tall.

Now consider a case in which the degree of Feynman's tallness is well-known, but the prevailing standard of tallness is uncertain. This is where you ask me what counts as tall in Los Alamos, and I reply by saying "Around here, ..." and continue by uttering (6).

(8)	c	$\mathbf{d}(c)(\llbracket tall \rrbracket)$	$\iota(\max(\lambda d.\mathbf{tall}(d, \mathbf{f})))$	$c \in \mathbf{tall}(\mathbf{d}(c)(\llbracket tall \rrbracket), \mathbf{f})$
	c_5	184	181	no
	c_6	182	181	no
	c_7	180	181	yes
	c_8	178	181	yes

According to this second scenario, the discourse participants both know exactly how tall Feynman is. Thus the candidates under consideration do not differ as to the maximal degree of Feynman's tallness. However, there is considerable uncertainty about what counts as tall in this circumstance of utterance. According to the chart in (8), there are four degrees of tallness under consideration as the standard of absolute tallness, each one progressively less stringent than the last. Only the standards for c_7 and c_8 are sufficiently low standards that Feynman counts as tall. Therefore if c_5, \dots, c_8 is the context at the point in a discourse immediately before uttering (6), the updated context will contain only c_7 and c_8 : $\llbracket Feynman is tall \rrbracket(\{c_5, c_6, c_7, c_8\}) = \{c_7, c_8\}$. We have learned nothing new about Feynman's height by accepting (6), but we do learn something about the prevailing standard of tallness. In such a discourse situation, the update effect of (6) is purely sharpening.

Normally, of course, there is uncertainty both about what is the case and about the prevailing vague standards. We may know that Feynman has some imprecise but high degree of height, and

we may also suspect that the standard of tallness at Los Alamos might be higher than elsewhere. In such a state of knowledge all eight of the candidates considered in the last two scenarios might be live alternatives. Update with the denotation of (6) would narrow the set of candidates to c_1 , c_2 , c_7 , and c_8 : $\llbracket \text{Feynman is tall} \rrbracket (\{c_1, \dots, c_8\}) = \{c_1, c_2, c_7, c_8\}$. We would have simultaneously refined our knowledge of the degree to which Feynman is tall, and also our knowledge of how the word *tall* is used at Los Alamos. That is, in the normal situation, an utterance of a sentence like (6) simultaneously operates at both a descriptive and a metalinguistic level.

The somewhat idealized but familiar picture of discourse that emerges is a progressive refinement of the assumptions shared by the interlocutors. But are candidate standards really discarded for good, or only temporarily? After all, if I tell you *Give this tall glass to that short basketball player*, you do not object to my description of the basketball player on the grounds that her height exceeds any standard of tallness that could justify calling a drinking glass tall. Assuming that being short entails being not tall, it seems that I must have reset my standard of tallness in between uttering the direct object and the indirect object. But this is just an example of the well-known (but poorly understood) dependence of vague standards on some relevant comparison class (see Klein (1991:685) for discussion and references): what matters is whether a glass is tall *for a glass*, or whether a person is short *for a basketball player*. To see that the sharpening entailed by the use of *tall* remains in full force, note that it would be inconsistent of me to continue by saying *And by the way, this even taller glass is not a tall glass, and that even shorter basketball player is not a short basketball player*. Thus the original utterance does commit me to certain persistent standards of tallness, though it does seem necessary to keep track of a different standard for each comparison class (see Kamp and Partee (1995) and Kennedy (1997) for discussion and proposals for how to do this).

3. Dynamic comparison and higher-order vagueness

Any complete semantic account of gradable adjectives must provide a treatment for comparatives. However, comparatives are quite complex both syntactically and semantically, and it is clearly not possible to give a comprehensive discussion here of the many analyses available in the literature (von Stechow (1984), Klein (1991), and Kennedy (1997) give useful surveys). Nevertheless, in order to give some concrete reason to suppose that the dynamic approach can lead to a reasonable account of comparatives, I will show how at least two well-known types of analyses of the comparative can be reconstructed in a dynamic framework. The main point of interest for the larger issues of this paper is that the truth conditions proposed here automatically predict that measure phrases and comparatives, unlike normal absolute uses, have no sharpening effect. In addition, the treatment of degree modifiers leads to an account of higher-order vagueness.

3.1. Measure phrases

Providing an analysis for measure phrases (e.g., *two meters*, as in *Mary is two meters tall*) will introduce the ingredients needed for an analysis of comparatives and degree modifiers.

It will be convenient to adopt a notational convention adapted from formal treatments of quantification. In many theories, a quantifier binds a pronoun by manipulating the contextually-supplied assignment function that would otherwise provide the pronoun with a deictic interpretation. Since

we will need to quantify over vague standards rather than over variable assignments, what we need is a modified delineation function: therefore let $c[d/\alpha]$ be a world exactly like c except perhaps as required in order to guarantee that $\mathbf{d}(c)(\alpha) = d$.

Using modified delineations, let the denotation of $\llbracket two\ meters \rrbracket$ be $\lambda\alpha\lambda x\lambda C.\{c \in C : c[2m/\alpha] \in \alpha(x)\}$, so that $\llbracket Mary\ is\ two\ meters\ tall \rrbracket = c[2m/\llbracket tall \rrbracket] \in \llbracket tall \rrbracket(\mathbf{m})$. Then the sentence *Mary is two meters tall* is true just in case Mary would count as tall if the standard for tallness were set to two meters.

It is important that $c[2m/\llbracket tall \rrbracket] \in \llbracket tall \rrbracket(\mathbf{m})$ can be true even if $c[2m/\llbracket tall \rrbracket]$ is not in the common ground, since a measure phrase can impose a temporary standard that is not remotely under consideration as a standard for absolute tallness.

Even more important, note that $c[2m/\llbracket tall \rrbracket] \in \llbracket tall \rrbracket(\mathbf{m})$ can be true even if $c \notin \llbracket tall \rrbracket(\mathbf{m})$ —that is, Mary can be six feet tall in c even if Mary does not qualify as tall in c . Because the truth conditions for measure phrases refer only to modified delineations, they eliminate worlds from the context based only on descriptive facts, without regard to vague standards. In other words, unlike absolute uses of gradable adjectives, measure phrases have only a descriptive mode of use, and no metalinguistic mode.

3.2. Comparatives

One of the fundamental issues in the theory of comparatives is whether the comparative or the absolute form is semantically more basic. In English (as in many other languages; see Klein (1980, 1991)), the comparative is morphologically or syntactically more complex than the absolute form of an adjective (*taller* = *tall* + *-er*, *more lucky* = *more* + *lucky*), suggesting that the meaning of the comparative ought to be a compositional function of the meaning of the absolute. Nevertheless, many semantic treatments work the other way around, basing the meaning of the absolute form on the meaning of the comparative. In contrast to such accounts, I will provide a semantics consistent with English morphology, i.e., I will express the meaning of the comparative as a function of the meaning of the absolute.

One standard approach to comparatives quantifies over degrees. Klein (1991:685) gives the provenance of the basic idea, and offers a defense against some criticisms. Moltmann (1992), Lerner and Pinkal (1995), and Gawron (1995) can be classified as similar. On this type of analysis, the denotation of the comparative suffix can be rendered as follows:

$$(9) \quad \llbracket -er \rrbracket = \lambda\alpha\lambda x\lambda y\lambda C.\{c \in C : \exists d.(c[d/\alpha] \in \alpha(y)) \wedge (\neg c[d/\alpha] \in \alpha(x))\}$$

This meaning quantifies over degrees: for each degree d , the truth conditions evaluate an expression based on the absolute meaning of the adjective with respect to a candidate in which the relevant standard has been temporarily set to d . Combining (9) with the adjectival denotations as described in section 2, we have:

- (10) a. Feynman is taller than Bill.
 b. $\lambda C.\{c \in C : \exists d.(c[d/\llbracket tall \rrbracket] \in \llbracket tall \rrbracket(\mathbf{f})) \wedge (\neg c[d/\llbracket tall \rrbracket] \in \llbracket tall \rrbracket(\mathbf{b}))\}$

The truth conditions in (10b) can be paraphrased as claiming that there is a degree d such that Feynman is at least that tall but Bill is not.

The main alternative type of analysis, which Klein associates with von Stechow (1984), as taken up by, e.g., Heim (1985) and Izvorski (1995), treats the comparative as involving definite descriptions of a degree. A definite-description type analysis can be couched in the current framework along the following lines:

- (11) a. $\llbracket -er \rrbracket = \lambda\alpha\lambda x\lambda y\lambda C.\{c \in C : \exists d > 0 : c[d + \iota(\max(\lambda d'.c[d'/\alpha] \in \alpha(x)))/\alpha] \in \alpha(y)\}$
 b. $\lambda C.\{c \in C : \exists d > 0 : c[d + \iota(\max(\lambda d'.c[d'/\llbracket tall \rrbracket] \in \llbracket tall \rrbracket(\mathbf{f})))/\llbracket tall \rrbracket] \in \llbracket tall \rrbracket(\mathbf{b})\}$

In (11a), $\iota(\max(\lambda d'.c[d'/\alpha] \in \alpha(x)))$ is the maximal degree to which x has the gradable property α in c ; for instance, if *Bill* is exactly two meters tall in c , $\iota(\max(\lambda d'.c[d'/\llbracket tall \rrbracket] \in \llbracket tall \rrbracket(\mathbf{b})))$ is two meters. Assuming we have a suitable sum operator ‘+’ defined over the set of degrees whose zero element is 0, (10) comes out as asserting that the maximal degree to which Feynman is tall exceeds the maximal degree to which Bill is tall by some positive amount d .

As far as I know, giving a dynamic treatment of gradable adjectives is neutral with respect to which style of analysis is preferable; see Klein (1991), Rullman (1995), Kennedy (1997, to appear), and many others for discussion. The main point here is that the dynamic treatment of gradable adjectives is compatible with at least two respectable theories of the meaning of the comparative.

In any case, it is clear that just as for measure phrases, the semantics for comparatives here guarantee that a use of a comparative will have no sharpening update effect whatsoever (at least, no sharpening of the adjective under comparison).

3.3. Degree modifiers

Let **very** be a relation over degrees such that **very**(s, d, d') holds only if the difference between d and d' is larger than the relevant vague standard s .

- (12) $\llbracket very \rrbracket = \lambda\alpha\lambda x\lambda C.\{c \in \alpha(x)(C) : \exists d.(c[d/\alpha] \in \alpha(x)(C)) \wedge \mathbf{very}(\mathbf{d}(c)(\llbracket very \rrbracket), \mathbf{d}(c)(\alpha), d)\}$

This denotation for *very* is essentially a dynamic version of the analysis of Kennedy and McNally (1999). Here, $\llbracket very \rrbracket$ is an adjective modifier: it takes an adjective meaning (α) and returns an adjective meaning (a function from individuals x to an update function). *Feynman is very tall* will be true just in those worlds c in which Feynman is tall (i.e., $c \in \alpha(x)(C)$), and in which the maximal degree of Feynman’s tallness exceeds the general standard of tallness in c by an amount sufficient to qualify as **very**. That is, Feynman must be tall and then some.

Of course, *very* itself is vague, and so its update effect depends on the prevailing standard $\mathbf{d}(c)(\llbracket very \rrbracket)$. Thus update with *very* will normally also sharpen knowledge of the vague standard associated with *very*. For instance, if the common ground provides a fairly precise notion of how tall is tall, and a precise notion of how tall Feynman is, a use of *Feynman is very tall* will reveal in part how much taller than tall someone must be to count as very tall.

This analysis allows stacking of degree modifiers in such a way that *very very tall* results in

more stringent truth conditions than *very tall*. The denotation of $\llbracket \textit{very tall} \rrbracket$ is the semantic type of a gradable adjective, which means that Feynman will be very, very tall just in case his tallness exceeds by a sufficient amount a degree that itself qualifies as very tall.

Here is an example of the sort of predictions this theory makes. Since length and height are commensurate (which is why it is possible to say *this bridge is longer than it is tall*), it makes sense to assume that *long* and *tall* are true of the same set of degrees. In a world in which the standard of tallness is equal to the standard of length (i.e., for c such that $\mathbf{d}(c)(\llbracket \textit{tall} \rrbracket) = \mathbf{d}(c)(\llbracket \textit{long} \rrbracket)$), then in that same world any object that qualifies as very tall ought to also qualify as very long (subject to appropriate orientation with respect to the ground and/or the observer). As near as I can judge, this prediction seems to be true.

3.4. Higher-order vagueness as vagueness quantifiers

One of the advantages of the dynamic treatment of vagueness is that it leads to a natural treatment of higher-order vagueness.

Williamson (1994:chapter 5) criticizes some supervaluation theories of vagueness in part on the grounds that they break down when confronted with higher-order vagueness. The basic problem is that simple supervaluation theories recognize only three crisp possibilities: a vague predicate clearly applies, it clearly fails to apply, or it neither applies nor fails to apply. The danger, as Fine (1975:297) puts it, is that on supervaluation treatments “anything that smacks of being a borderline case is treated as a clear borderline case”. But just as there can be uncertainty about precisely where the border between tall and not tall falls, there can be second-order uncertainty about the boundary between clearly tall and borderline tall: is Feynman clearly clearly tall, or only sort of clearly tall? An adequate theory must deal appropriately with second-order and higher orders of vagueness.

Fine (1975, section 5) and Williamson (1998) elaborate supervaluation theories to handle a *definitely* operator. Williamson adds a modal-like notion of accessibility over the class of delineations, so that *Feynman is definitely tall* comes out as true at a world c just in case *Feynman is tall* is true in c under every delineation accessible from c . As long as the accessibility relation R is not both transitive and symmetric, *definitely definitely p* can receive a value different from *definitely p*, and we have higher orders of vagueness as desired. The problem then becomes finding and justifying a suitable choice for the accessibility relation.

On the dynamic approach, independently-motivated details of the semantics gives rise to an account of higher-order vagueness automatically, without stipulating an accessibility relation over delineations. The basic idea is to allow modifiers such as *definitely*, *indisputably*, *undoubtedly*, *unquestionably*, *clearly*, *sort of*, *arguably*, etc., to quantify over worlds in a context, very much in the same way that nominal or adverbial quantifiers quantify over individuals or situations. That is, I will propose that modifiers like *clearly* are VAGUENESS QUANTIFIERS. Thus unlike being very tall, being definitely tall cannot be determined by looking at a single candidate. Instead, it is a property whose applicability depends both on a specific candidate and on the context as a whole. More specifically, Feynman is definitely tall in $c \in C$ just in case he is tall in c and there is no other candidate $c' \in C$ in which Feynman is just as tall but does not exceed the standard of tallness for c' .

$$(13) \quad \llbracket \textit{definitely} \rrbracket = \lambda\alpha\lambda x\lambda C.\{c \in \alpha(x)(C) : \forall d.(c[d/\alpha] \in C) \rightarrow c[d/\alpha] \in \alpha(x)(C)\}$$

(14)	c	$\mathbf{d}(c)(\llbracket \textit{tall} \rrbracket)$	$\iota(\max(\lambda d.\mathbf{tall}(d, \mathbf{f})))$	$c \in \llbracket \textit{tall} \rrbracket(\mathbf{f})$	$c \in \llbracket \textit{definitely tall} \rrbracket(\mathbf{f})$
	c_1	180	181	yes	no
	c_2	182	181	no	no
	c_3	180	183	yes	yes
	c_4	182	184	yes	yes

In this context there is some uncertainty about the standard of tallness: if you are shorter than 180 centimeters, you are definitely not tall; if you are between 180 and 182 centimeters, you might be tall; and if you are taller than 182 centimeters, you are definitely tall. Candidate c_1 survives update with *Feynman is tall* because c_1 is a candidate in which the standard of tallness is 180 centimeters, and Feynman is taller than that in c_1 ; c_1 does not survive update with *Feynman is definitely tall*, however, because although Feynman is tall in c_1 , there is another candidate—namely, c_2 —in which Feynman is just as tall, but in which he nevertheless does not count as tall because the standard of tallness is set higher in c_2 (i.e., $c_1[182/\llbracket \textit{tall} \rrbracket] = c_2$). Both candidates c_3 and c_4 survive update with *Feynman is definitely tall* because in those candidates Feynman’s tallness is beyond the haze of vague uncertainty.

On the dynamic approach, then, *Feynman is definitely tall* comes out as a statement about the condition of the current discourse, roughly: ‘the range of delineations currently under consideration all support the claim that someone of Feynman’s height counts as tall’.

The denotation for *definitely* according to (13) is not a gradable property, and so quite unrealistically does not recognize that *definitely* itself is vague. As a result, it is a theorem of the system that *definitely definitely tall* means the same thing as *definitely tall*, contrary to most people’s intuitions (though such truth conditions seem to be appropriate for some people for examples like *indisputably indisputably tall*). This would mean that we have first-order vagueness but no higher forms of vagueness. Rather than revising the definition of *definitely*, it will be convenient to provide instead an explicitly gradable denotation for *clearly*, on the understanding that the analysis of *definitely* will probably have to be complicated in a similar fashion. The denotation of *clearly* in (15) is a combination of elements from the denotations given above in (12) and (13) for *very* and *definitely*:

$$(15) \quad \lambda\alpha\lambda x\lambda C.\{c \in \alpha(x)(C) : \mathbf{clearly}(\mathbf{d}(c)(\llbracket \textit{clearly} \rrbracket), \\ \iota(\max(\lambda d.c[d/\alpha] \in \alpha(x))), \\ \iota(\max(\lambda d.c[d/\alpha] \in C)))\}$$

Not surprisingly, making a first-order modifier vague will lead directly to higher orders of vagueness. In order for $\llbracket \textit{clearly tall} \rrbracket$ to apply to Feynman at a world c , the maximal degree of Feynman’s height in c (which is given by $\iota(\max(\lambda d.c[d/\llbracket \textit{tall} \rrbracket] \in \llbracket \textit{tall} \rrbracket(\mathbf{f})))$) must exceed the highest standard

for height under consideration for worlds like c (given by $\iota(\max(\lambda d.c[d/[[tall]]] \in C))$) at least by a certain vague amount (namely, $\mathbf{d}(c)([[clearly]])$). In other words, in order to be clearly tall Feynman must be definitely tall (in the sense of (13)) and then some. For instance, given a context as depicted in (14), assuming that $\mathbf{d}(c)([[clearly]]) = 0.6\text{cm}$ for $c = c_1, c_2, c_3, c_4$, then c_3 and c_4 survive update with $[[clearly tall]]$, since the maximal height of Feynman is more than 0.6 centimeters greater than the highest standard of tallness under consideration. But only c_4 survives update with $[[clearly [clearly tall]]]$, because only in c_4 is Feynman at least 0.6 cm taller than he needs to be to qualify as clearly tall.⁶

For expressions exhibiting only first-order vagueness, truth at a world depends only on that world (and the delineation that is part of that world). For expressions leading to second-order vagueness, truth at a world depends on that world and also on the delineations associated with other worlds in the evaluation context. Second-order and higher-order vagueness arises from the interaction of recursive context update in combination with the recognition that operators such as *clearly* are themselves vague.

The proposal is that being clearly tall means being tall enough to count as tall in every contextually relevant world. Because the dynamic semantics makes the entire set of contextually relevant alternatives available to the denotation of an expression, it is possible for the denotation of a vagueness quantifier like *clearly* to quantify over the context set. Predicates that give rise to higher-order vagueness merely exploit contextual information in the same way that nominal or adverbial quantifiers do.

The picture that emerges is one on which measure phrases and comparatives do not have any sharpening potential, but degree modifiers like *very* and *clearly* do. In each case, the sharpening potential follows directly from stating the truth conditions of the expression in question within a dynamic framework.

4. Relative stupidity

The second case study investigates a particular subclass of those adjectives that take infinitival complements that I will call ‘stupid’ adjectives. I will argue that stupid adjectives are in one sense the opposite of comparatives: whereas the truth conditions of comparatives guarantee that they will have descriptive entailments but no sharpening effect, those of stupid adjectives will have sharpening entailments, but no descriptive effect.

I will use *stupid*, *smart*, and *lucky* as my prototypical examples of three slightly different sub-types of stupid adjective. Relative uses of stupid adjectives are quite common, as well as utterly natural and idiomatic across a wide range of registers. A few naturally-occurring examples will give an impression:⁷

⁶Some justification for this constituent structure is in order. Since *clearly* can certainly modify adjectives, and since *clearly tall* is a (complex) adjective, the structure in the text ($[[clearly [clearly tall]]]$) must be available. Some linguists have a strong intuition that $[[clearly clearly] tall]$ is also a possible compositional structure. I’m not aware of any obstacles to extending the analysis to the alternative structure; in any case (though the relevant judgments are somewhat delicate), as near as I can tell the truth conditions for the two structures would have to be equivalent. More philosophically-minded discussions insist on $[[clearly [clearly \dots]]]$ as the compositional structure, which is one reason I use it here.

- (16) a. Steel managers were **stupid** to still be building open-hearth furnaces while others had switched to oxygen furnaces.
- b. A game called 'rck.video' innocently entertains the user with an animation of Madonna before surreptitiously erasing all the computer's files. Then it gloats, "You're **stupid** to download a video about rock stars."
- c. Still, GE and Whirlpool are **smart** to tout their replacement offers at a time when many people dread dealing with repairs.
- d. In hindsight, some companies were **smart** to sell stock before the market debacle.
- e. You did well to keep the hotel calm, and you were **lucky** to find the documents the way you did.

Some other predicates that seem to behave just like *stupid*, *smart* and *lucky* in the relevant respects include the following, taken from lists due to Rivière (1983) and Wilkinson (1970:425, 1976): *brave*, *careless*, *clever*, *(in)considerate*, *courageous*, *cowardly*, *cruel*, *foolish*, *impudent*, *(un)kind*, *naughty*, *nice*, *(im)polite*, *right*, *rude*, *silly*, *stupid*, *wicked*, *(un)wise*, and *wrong*, to which I add here *crazy*, *dumb*, *evil*, *(un)lucky*, *mean*, *(im)prudent*, *smart*, and perhaps also *within one's rights*, *well-advised*, and *morally derelict*; there may be others.

The main point of this case study does not depend on proving that stupid adjectives constitute a natural class that is homogeneous either syntactically or semantically. Rather, if even one adjective (say, *stupid* or *lucky*) has the crucial semantic properties when it occurs with an infinitival complement (namely, no descriptive entailments beyond presuppositions and sharpening), that would be sufficient. Nevertheless, to the extent that there are many distinct adjectives that have the relevant properties (as I believe there are), and to the extent that there are a number of distinct constructions with analogous properties (see section 4.5), then so much the better for our confidence that we are dealing with a robust phenomenon.

The next two sections make the case that (17) has no descriptive entailments.

- (17) Feynman was stupid to dance like that.

In relative uses like (17), *stupid* takes two arguments, an infinitival VP and a subject NP. Thus the claim at hand divides into two main components: the nature of the entailments associated with the infinitival complement (section 4.1), and those associated with the subject position (section 4.2).

Sections 4.3, 4.4, and 4.5 provide a formal analysis that leads to an explanation for the unembeddability phenomenon mentioned in the introduction (**Feynman wanted to be stupid to dance like that*). Section 4.6 discusses the interaction of stupid adjectives with counterfactual modality (*Feynman would be stupid to dance like that*).

⁷(16a)–(16d) are from Wall Street Journal articles written (respectively) by Lester C. Thurow, 12 June 1987; Ronald Alsop, 25 June 1987; Linda Sandler, 4 November 1987; and Asra Q. Nomani, 17 June 1988. Example (16e) is from Martin Cruz Smith (1999) *Havana Bay*, Random House, New York, p. 183.

4.1. VP-oriented entailments: Stupid adjectives presuppose their VP complement

Stupid adjectives do have entailments with respect to their VP complements, but those entailments are presupposed and not asserted.

- (18) a. Feynman was stupid to dance on the table.
b. Feynman wasn't stupid to dance on the table.

Wilkinson (1970) notes that sentences like (18a) and their negation in (18b) both entail that Feynman danced. An entailment that remains constant under negation, of course, is the hallmark of presupposition.

- (19) a. Was(n't) Feynman stupid to dance on the table? yes/no Q
b. Who was stupid to dance on the table? WH Q
c. Perhaps Feynman was stupid to dance on the table. possibility
d. If Feynman was stupid to dance on the table, then tell him. conditional

In addition, (19a) through (19d) show that the entailment projects to the containing S under yes/no question formation, WH question formation, in the presence of epistemic modality, and when embedded in the antecedent of a conditional, exactly as expected if the content of the infinitival VP is presupposed rather than asserted.

Another argument that we are dealing with presupposition comes from existential commitment:

- (20) a. Feynman was(n't) likely/eager/ready to talk to a student about it.
b. Feynman was(n't) stupid/smart/lucky to talk to a student about it.

The normal adjectives in (20a) do not necessarily entail the existence of any specific student. The stupid adjectives in (20b), however, do, whether negation is present or not. This is as we would expect if the infinitival complement of *stupid* is presupposed, since that would entail that the speaking event occurred and therefore that there is some student that Feynman actually talked to. Thus only (20b) can always felicitously be followed by anaphoric references to the student, such as *She immediately called the student newspaper*. Such a continuation is appropriate after (20a) only if the indefinite *a student* receives a specific interpretation, or as an instance of modal subordination (e.g., *She would have immediately called the newspaper*).

Whatever you think presuppositions are, relative uses of stupid adjectives undoubtedly have them with respect to the denotation of their infinitival VPs. (See section 4.6 for an discussion of apparent counterexamples involving counterfactual *would*.)

4.2. Subject-oriented entailments: various volitionality presuppositions

Determining whether stupid adjectives associate descriptive entailments with (the referent of) their subjects requires a close look at entailments concerning sentience and volitionality.

Some predicates (e.g., raising predicates such as *likely*) do not even have an external argument (i.e., they do not select for a subject), and so obviously cannot have any subject-specific descriptive entailments. Both Rivière (1983) and Wilkinson (1970) mistakenly classify stupid adjectives as raising predicates, citing the fact that stupid adjectives can take a sentential subject (*For Feynman to dance like that is likely/*eager/stupid*) and that the sentential subject can be extraposed (*It is likely/*eager/stupid for Feynman to dance like that*). But, as Wilkinson later realized (see Wilkinson 1976), it is clear that stupid adjectives do select for a subject (take an external argument), since they do not accept pleonastics or idiom chunks as subjects:

- (21) a. *It is stupid to be raining.
b. *There is stupid to be a party tomorrow.
c. *Tabs were lucky to be kept on her.

More to the point here, stupid adjectives uniformly entail that the referent of the subject is sentient (i.e., capable of intentional action).

- (22) a. It was stupid for the carpet to be cleaned right before the big party.
b. #The carpet was stupid to be cleaned right before the big party.

There is a clear intuition that the contrast between (22a) and (22b) is due to the implausible implication that the carpet must be sentient in (22b). It is easy to prove that this implication projects through negation, modal possibility, question formation, and so on—that is, it is a presupposition. And because it is a presupposition, this sentience requirement is consistent with the main claim that relative uses of stupid adjectives convey no descriptive entailments apart from sharpening.

It is important here to distinguish between sentience on the one hand and volitional involvement in the described event on the other. Sentience is a capability, not a condition, so it is perfectly possible to presuppose sentience without entailing anything specific about the mental state of the subject participant during the eventuality described by the event. For instance, the predicate *unconscious* presupposes sentience, but does not entail volitional involvement in the unconscious state. By the same token, relative uses of *lucky* presuppose sentience (**The carpet was lucky to be cleaned*), but do not require any volitional involvement in the eventuality described by the infinitival:

- (23) Feynman was lucky to be born in the 20th century.

This sentence says nothing specific about Feynman's mental life, either when he was born or at any point later in his life.

However, some stupid adjectives entail in addition to sentience that the subject must be volitionally involved in the described eventuality.

- (24) ?Feynman was smart to constrict his pupils when the intense light hit them.

This sentence is odd because *smart* entails that Feynman had discretion over his pupillary reflexes, which is normally physiologically impossible.

This type of volitional involvement is strongly reminiscent of Farkas' (1988:36) notion of RESPONSIBILITY: an individual *i* and a situation *s* are in the RESP relation just in case “*i* brings *s* about, i.e., just in case *s* is the result of some act performed by *s* with the intention of bringing *s* about.” Farkas claims that the RESP relation holds between arguments of some, but not all (p. 52) control predicates, and I will assume that RESP applies to some but not all stupid predicates. Among those for which it holds are *smart, clever, nice, right, etc.*

However, RESP does not apply to the polar counterparts of these adjectives (*stupid, dumb, mean, and wrong, respectively*). Rivière cites *John was stupid to fall into that ditch*: certainly it is not necessary for John to have acted with the intention of bringing about a falling. Rather, John must indeed have had some intention, but that intention must have been defective in failing to consider all of the consequences of his actions, since one of those consequences lead to his falling. Rivière suggests that the subject must be able “to choose between a process and its negation”, i.e., John must have been able to avoid falling in the ditch. Yet that is not quite right either.

(25) ??You were stupid to have twins.

Presumably it was possible to choose between the process that leads to twins versus its negation (having no children). The condition imposed by *stupid* is slightly stronger: it must be within the power of the subject participant to bring about the state of affairs described by the infinitival. Thus what is lacking in (25) is the ability (under normal circumstances) to effectively and reliably choose whether a pregnancy will involve twins rather than some other number of children.

These considerations distinguish three subtypes of stupid adjective:

(26)		<i>lucky</i>	<i>stupid</i>	<i>smart</i>	
	SENTIENCE	The subject must be capable of volition	yes	yes	yes
	DISCRETION	It is within the power of the subject to choose to bring about the situation described by the infinitive	neutral	yes	yes
	INTENTIONALITY	The subject intends for the situation described by the infinitive to come about	neutral	neutral	yes

The conjunction of discretion and intentionality is equivalent to Farkas' RESP.

Doubtless there are other subtleties and refinements. For instance, it may be that *lucky* also requires the negation of discretion, or that at least some significant aspect of the infinitival situation be out of the control of the subject. However, for present purposes, it does not matter what the precise formulation of these various conditions are. The reason is that they are all presuppositions, and therefore consistent with the main claim that stupid adjectives have no descriptive entailments. Not only do any implications of responsibility and discretion project when embedded under negation, question formation, conditionals, etc. (i.e., the standard tests for presuppositions as in (18) and (19)), sentences that describe situations in which these implications fail to hold such as (24) and (25) are infelicitous rather than merely false.

Even assuming that the requirements as given in (26) are presuppositions, it is still conceivable

that stupid adjectives impose descriptive entailments above and beyond these requirements. For comparison, consider *proud*, which, like stupid adjectives, presupposes its infinitival complement and presupposes that its subject is sentient, but descriptively entails in addition that its subject participant has a certain emotional state.

Crucially, unlike *proud* or *glad*, stupid adjectives fail to entail anything specific about the emotional or mental state of their subject. Speakers are quite confident in making judgments such as *Feynman was stupid to dance like that* without knowing anything about the subject's mental state. Feynman's dancing can qualify as stupid no matter what Feynman's mood or intentions: at any point before, during, or after the event he can be angry, happy, ashamed, worried, or blithely unaware of the consequences of his actions.

In other words, the applicability of a relative use of a stupid adjective depends only on externally observable criteria. What matters is the outcome, and ignorance is no excuse:

(27) John was rude to address the Queen.

Despite the best of intentions, due to John's ignorance of British royal protocol (which specifies that one must never initiate a conversation with the Queen), his behavior unquestionably counts as rude.

To sum up, when sentience, discretion, and intentionality entailments are present for stupid adjectives, they are presuppositions. Several stupid adjectives, including *lucky* (also *unlucky*, *fortunate*, and perhaps some others), have no presuppositions beyond sentience, and so the *lucky* subclass constitutes an especially clear and straightforward example of an adjective that has no entailments beyond its presuppositions and its sharpening effect.

4.3. Infinitival denotations

Extending the analysis in section 2 to stupid adjectives requires first considering the meaning of the infinitival VP. Portner (1992), building on Bresnan (1972), argues that such infinitives are essentially futurate: they denote a set of situations, each of which includes the current situation as an initial subpart and extends into a possible future at least far enough to also include a subsituation in which the content of the infinitive obtains.

Crucially, such situations may not be large enough to qualify as complete worlds. Therefore, it is necessary to further articulate the ontology. I will assume a situation-based semantics roughly as proposed in Kratzer (1989) and developed in Portner (1992). There will be a set U of individuals, a set S of situations such that $U \subset S$, and a partial order \leq defined on S , the part-of relation. Possible worlds are situations that are not proper parts of larger situations: a situation s is a complete world iff $\forall s' \in S[(s \leq s') \rightarrow (s = s')]$. Contexts continue to contain only complete worlds.

Thus if I say

(28) I am happy to give you this award. [Portner]

I am happy in relation to a set of situations, each of which begins with the utterance situation and

extends into the future to encompass a giving event (see Portner (1992:37) for remarks about the temporal properties of situations). There may be several such situations, each corresponding to a different way in which the future may unfold.

It is important that the situations in the denotation of the infinitival need not be entire possible worlds, but can be smaller situations.

(29) [Feynman dances wildly and spills beer on his shirt.]

Me: Feynman is stupid to dance like that.

You: No, Feynman **ISn't** stupid to dance like that. What you don't realize is that he is distracting that spy over there with his antics while the FBI closes in.

Your denial does not call into question what happened, and you do not deny that Feynman danced in the manner I attribute to him. That is presupposed, and you do not challenge that presupposition. In fact, you don't even deny that dancing wildly enough to spill beer on oneself meets the standard of stupidity under normal circumstances. Rather, what you are doing is encouraging me to take a wider view of the situation in order to direct my attention to extenuating circumstances.

I will take the phrase “extenuating circumstances” rather literally. What the objector in (29) is doing is insisting on extending each situation in the extension of the VP so as to include a larger chunk of the world: not only the dancing and its immediate consequences, but some longer-range objectives as well.

Let's assume that when I uttered my claim in (29), I had in mind for the denotation of the infinitival VP a set of Portnerian situations that included Feynman, dancing, beer, Feynman's shirt, and nothing else. Participating in that restricted set of situations does indeed render Feynman relatively stupid. Your challenge goes to my selection of a set of relevant situations. You would like to suggest a set of situations that contains more. For each situation s in the denotation of my VP, you nominate another situation s' such that $s \leq s'$, where s' contains not only Feynman and the beer but also a spy and some FBI agents. Relative to that set of situations, Feynman may not be stupid—he may in fact be preventing the nuclear secrets of his nation from falling into enemy hands.

4.4. *Presuppositions and relative stupidity*

One final piece of technical machinery before providing an analysis of relative stupidity: how presuppositions work. As explained e.g., by Heim (1982, 1983) or Beaver (1997, to appear), presuppositions can be treated as constraints on the input context: update will proceed only if the context satisfies the relevant presupposition. More technically, update functions are rendered as partial functions that are only defined on contexts that satisfy all relevant presuppositions.

(30) $\forall P, x, C : \llbracket \textit{stupid}_{\text{rel}} \rrbracket(P)(x)(C)$ is defined only if $C \subseteq [\text{SEN}(x) \wedge \text{DISC}(x, P(x)) \wedge P(x)]$,

in which case $\llbracket \textit{stupid}_{\text{rel}} \rrbracket = \lambda P \lambda x \lambda C. \{c \in C : c \in \text{stupid}(\mathbf{d}(c)(\llbracket \textit{stupid} \rrbracket), x, P)\}$.

According to (30), the denotation of a relative use of *stupid* isn't even defined unless three presuppositions are met: the subject must be sentient, the subject must have discretionary power over the state

of affairs described by the infinitival, and the proposition formed by applying the infinitival property P to the individual x must be entailed by the input context. In the style of dynamic update, these requirements are expressed by stating that the evaluation context must be a subset of the set of worlds in which these conditions obtain. If the presuppositions hold, then the update effect of the relative use of the adjective is to filter out those worlds in which the standard of stupidity ($\mathbf{d}(c)(\llbracket stupid \rrbracket)$) is too high for x 's participation in the event $P(x)$ to count as stupid.

Here, then, is the update effect of *Feynman is stupid to dance*:

$$(31) \quad \text{CCP:} \quad \lambda C. \{c \in C : c \in \mathbf{stupid}(\mathbf{d}(c)(\llbracket stupid \rrbracket), \mathbf{f}, \llbracket to\ dance \rrbracket)\}$$

$$\text{Presup.:} \quad \forall C, \llbracket Feynman\ is\ stupid\ to\ dance \rrbracket(C) \text{ is defined iff}$$

$$C \subseteq [\text{SEN}(\mathbf{f}) \wedge \text{DISC}(\mathbf{f}, \llbracket to\ dance \rrbracket(\mathbf{f})) \wedge \llbracket to\ dance \rrbracket(\mathbf{f})]$$

A use of *Feynman is stupid to dance* presupposes that for every candidate c in the initial context, Feynman danced in c . Therefore a situation in which Feynman may not have danced leads to presupposition failure:

$$(32) \quad c \quad \mathbf{d}(c)(\llbracket stupid \rrbracket) \quad c \in \llbracket to\ dance \rrbracket(\mathbf{f}) \quad c \in \mathbf{stupid}(\mathbf{d}(c)(\llbracket stupid \rrbracket),$$

$$\mathbf{f},$$

$$\llbracket to\ dance \rrbracket)$$

c_9	d_9	no	#
c_{10}	d_{10}	no	#
c_{11}	d_9	yes	no
c_{12}	d_{10}	yes	yes

Feynman does not dance in c_9 or c_{10} , and this means that the inclusion of either candidate in a context renders a use of (31) infelicitous, as indicated by the ‘#’ in the rightmost column. Thus $\llbracket Feynman\ is\ stupid\ to\ dance \rrbracket(\{c_9, c_{10}, c_{11}, c_{12}\})$ is undefined.

Therefore assume that the input context contains only c_{11} and c_{12} . These two candidates differ only with respect to the standard of relative stupidity. It is now possible to deduce that $\mathbf{d}(c_{11})(\llbracket stupid \rrbracket) = d_9$ is greater than $\mathbf{d}(c_{12})(\llbracket stupid \rrbracket) = d_{10}$, since the same dancing situation counts as stupid with respect to d_{10} , but not with respect to d_9 . Thus the update effect of (31) will be to exclude candidate c_{11} , and so $\llbracket Feynman\ is\ stupid\ to\ dance\ like\ that \rrbracket(\{c_{11}, c_{12}\}) = \{c_{12}\}$. Given a context that satisfies the relevant presupposition, the only update effect of evaluating (31) is to exclude candidates in which the standard of relative stupidity is set too high.

In other words, the lexical semantics of stupid adjectives guarantees that their update effect involves sharpening and nothing else.

4.5. Explaining unembeddability: a prohibition against covert metalinguistic wishing

Sections 4.1 and 4.2 argued that when stupid adjectives take infinitival complements, they have no

descriptive entailments beyond presuppositions and sharpening entailments. Sections 4.3 and 4.4 showed how to express this hypothesis in a dynamic framework. This section shows how doing so leads to an explanation for unembeddability: as mentioned in the introduction, stupid adjectives cannot be embedded beneath control predicates when the stupid adjectives occur with an infinitival complement. A few additional examples:

- (33) a. I hoped to be [lucky (*to buy this lottery ticket)].
b. I intended to be [rude (*to be so honest)].
c. I preferred to be [wise (*to remain silent)].
d. I tried to be [fortunate (*to meet her)].

In each case, adding the infinitival complement in parentheses renders the sentence ungrammatical.⁸

The explanation for unembeddability will depend on assuming that a use of a control predicate is felicitous only in certain kinds of contexts: specifically, the content of its infinitival complement must not be entailed. Put informally, you can't want what you already have. Other authors have discussed this aspect of the semantics of control predicates before, specifically with respect to the predicate *want*. Heim (1992:194) observes "When a sentence like *I want you to call me on Monday* is used, there typically are doxastic alternatives where you do call me on Monday as well as those where you don't." Portner (1992:201) strengthens Heim's 'typically' to 'must': "Intuitively, wants must be for states of affairs that are believed to be as of yet undetermined as to whether they will be actualized". I will therefore call this the UNACTUALITY requirement, and I will go beyond Portner's discussion and attribute the unactuality requirement to control predicates in general rather than just to *want*.

The unactuality requirement resembles a presupposition in placing conditions on what can serve as an admissible evaluation context. Strictly speaking, however, unactuality is not a presupposition. It is not reducible to a condition on the individual worlds in the context, but rather constrains the context as a whole. In some sense, it is an anti-presupposition, since it is equivalent to a requirement that the infinitival not be presupposed. For example, consider how it correctly predicts that the following sentence (unrelated to adjectival constructions) is infelicitous:

⁸The longer versions of (33) are grammatical if the infinitival is interpreted as a purpose clause, in which case it would not be a complement of the adjective. (See Cormack (1998, section 4.1) for a general discussion of the complementation properties of various adjectives, including some discussion of stupid adjectives.) Other situations in which the infinitival does not function as a true complement must also be excluded:

- (i) a. Feynman was stupid enough to touch that light bulb.
b. Feynman was tall/boyish enough to touch that light bulb.
c. Feynman wanted to be stupid enough to touch that light bulb.

Wilkinson (1976) argues persuasively that *enough* augments the normal argument structure of the adjective it modifies by the addition of an infinitival VP. If so, then the infinitival VP in (ia) is not the infinitival VP subcategorized for by *stupid*. One reason to believe this is that *tall* or *boyish*, which do not normally subcategorize for an infinitival complement, can occur in the same frame, as shown in (ib). Because the *enough* construction has non-trivial descriptive entailments, it ought to be embeddable, and the perfect acceptability of (ic) bears out this prediction.

(34) #I want to realize that you are here.

The embedded verb *realize* is factive, so a use of (34) commits the speaker to the truth of the proposition that the listener is present. But if the speaker already assumes that the listener is present, then the speaker already realizes that the listener is present, and the content of the infinitival VP is guaranteed to be actual. Since it is infelicitous to want what you already believe to be the case, (34) is infelicitous by virtue of violating the unactuality requirement. Portner (1992) uses this type of semantic explanation extensively in his account of various impossible combinations of predicates and complement types (e.g., **They want that he is here*, **They wish that he will be here*, etc.).

The unactuality requirement forces any attempt to embed a relative use of a stupid adjective to have an unacceptable meaning.

(35) **Feynman wanted to be stupid to dance.*

Because the relative use of *stupid* presupposes its infinitival complement and has no entailments with respect to its subject position, the only way Feynman could be uncertain as to whether he was stupid to dance or not would be if the standard of stupidity were unsharp enough that the same dancing event qualifies as stupid in some candidate worlds but not in others. Thus in order for the content of Feynman's desire to become true, the standard of stupidity would have to be reset somehow. By hypothesis, the standard of stupidity is adjusted (sharpened) as discourse develops only as a side-effect of an assertion. Therefore Feynman's wish must be for the discourse in which (35) occurs to come to be in a different state; put another way, he is wishing for the word *stupid* to come to be used in a different way. Even assuming that Feynman was aware of the relevant conversation, and that it was indeed his wish that that discourse sharpen its standard of vagueness, it simply isn't possible to express such a wish in the form given in (35). To express such a metalinguistic desire, the linguistic elements would have to be part of the desire report, e.g., *Feynman wanted us to accept an assertion of "Feynman was stupid to dance"*.

The prohibition against covert metalinguistic wishing is not specific to vagueness or sharpening, but holds also for other potential interactions between a desire report and metalinguistic facts. The sentence *I want you to be me* cannot express my desire that you begin to speak in order for the second person pronoun to refer to me. Such a thought can only be expressed by putting the metalinguistic element into the content of the wish, as in *I want "you" to be me*, i.e., 'I want the word *you* to refer to me'. Similarly, the desire report *I want that bachelor to be married* cannot mean that I want you to agree that the word *married* should come to apply to a bachelor.

The claim is that stupid adjectives with infinitival complements are unembeddable because they have no descriptive entailments. If this hypothesis is correct, it predicts that whenever adjectives do have non-trivial descriptive entailments, they ought to be embeddable, and this prediction is borne out.

(36)		Presupposes all subject entailments?	Presupposes infinitival?	Embeddable?
a.	eager	no	no	I hope to be eager to go
b.	proud	no	yes	I expect to be proud to be your friend
c.	likely	yes	no	I intend to be likely to win
d.	stupid	yes	yes	*I want to be stupid to sing

The control predicate *eager* does not presuppose its infinitival, and it has distinct descriptive subject entailments, so it is definitely embeddable. Although *proud* does presuppose its complement, it also imputes an emotion to its subject. Consequently, as shown in (36b), *proud* is embeddable. Like all raising predicates, *likely* associates no entailments with its subject position; however, unlike stupid adjectives, it does not presuppose its infinitival VP, and that is enough to render it embeddable. It is only when a predicate both presupposes its infinitival and fails to associate any non-presupposed entailments with its subject—i.e., the characteristic semantic properties of stupid adjectives—that unembeddability results.

Could unembeddability possibly be explained in purely syntactic terms? This is unlikely, since there are other constructions that seem to have no descriptive entailments.

- (37) a. It was absurd/spiteful/proper of Feynman to dance like that.
 b. You did well/badly/all right to keep the hotel calm. (= (16e))

In particular, note that in both examples the infinitival is presupposed, and there are no descriptive entailments associated with the subject position (although (37b) shows the same kind of presupposition of volitional involvement discussed in section 4.2). And, just like relative uses of stupid adjectives, these constructions also cannot be embedded beneath a control predicate.

- (38) a. *I wanted it to be absurd of Feynman to dance like that.
 b. *I wanted you to do well to keep the hotel calm.

It seems unlikely that a syntactic explanation could naturally generalize over all three types of unembeddable construction, whereas it is easy to imagine extending the semantic explanation based on the interaction of the unactuality requirement with the dynamics of gradability.

4.6. Counterfactual stupidity

This section addresses a common objection to the analysis just given of stupid adjectives.

- (39) a. Feynman would be stupid to dance like that.
 b. You would be lucky to get a word in edgewise.

Although (39a) is consistent with a situation in which Feynman danced, it does not entail that he did; nor does (39b) guarantee that you get a word in edgewise. If stupid adjectives presuppose

the truth of their infinitival VP complement, why doesn't (39a) presuppose that Feynman danced? After all, it is well-known that at least some modal verbs are transparent for the purposes of presupposition projection.

- (40) a. Feynman found out that Ann is pregnant.
b. Perhaps Feynman found out that Ann is pregnant.
c. Feynman might find out that Ann is pregnant.
d. Ann is pregnant.

Since *find out* is factive, (40a) presupposes (40d). But since epistemic modality, including modal possibility, is transparent to presupposition projection, whether expressed as a sentence-level modifier (*perhaps*) or as a modal auxiliary (*might, may, could, etc.*), (40b) and (40c) also presuppose (40d). Thus if counterfactual *would* is the kind of modal that is transparent to presupposition projection, the facts in (39) are problematic.

Not to worry. Stupid adjectives are robustly factive by every standard test, including the modal possibility test. Recall that we saw an example of a relative use of a stupid adjective embedded under *might* in (19c), repeated here as (41a):

- (41) a. Perhaps Feynman was (not) stupid to dance like that.
b. Feynman might/may/could (not) have been stupid to dance like that.

Whether the possibility operator is sentence-level or a modal auxiliary, whether negated or not, the sentences in (41) clearly presuppose that Feynman danced.⁹

The key to understanding the sentences in (39), of course, is to realize that counterfactual conditionals are not reducible to epistemic modality. And in fact *would* does not in general allow presuppositions to project, as shown by these examples from Kasper (1992:314):

- (42) a. The king of France would be bald.
b. John would have managed to repair his car.
c. John would have stopped beating his dog as soon as it bit him.
d. John would have regretted having beaten his wife.

The sentences in (42) do not presuppose that France has a king, that John tried to repair his car, that John started beating his dog, or that John ever beat his wife.¹⁰

⁹Even if the presupposition did not project in (41), then stupid adjectives would be what Karttunen (1971) calls semi-factives (not to be confused with the morphological aspectual category 'semelfactive'): predicates that presuppose their complements under negation but not under possibility. For instance, according to Karttunen, *regret* is a proper factive, but *discover* is semi-factive: *He didn't discover that the world is round* entails that the world is round, but *He might discover that the world is round* allegedly does not. Even if stupid adjectives were semi-factive instead of factive, that would be sufficient for all of the main points in this paper to go through.

¹⁰In fact, it is tempting to decide that in each case they presuppose the negation of the main verb's presupposition; however, this is a misleading effect of considering them in a neutral context. To see that, e.g., (42c) is consistent with a context in which John had in fact been beating his dog, consider the following discourse: *I can tell you one reason why we should have*

Kasper argues that these *would* clauses function as the consequent of a counterfactual conditional with an unexpressed antecedent. For instance, Kasper renders the most natural construal of (42a) in a neutral context as *If there were a king of France, the King of France would be bald*. The corresponding paraphrase of (39a) would be *If Feynman had danced like that, he would have been stupid to dance like that*, which does not presuppose that the dancing occurred, in agreement with intuitions.

In order to correctly describe these intuitions, it is necessary to stipulate that bare *would* clauses like those in (42) must be evaluated with respect to some discourse-determined counterfactual context, rather than the common ground. Heim (1992) provides a good starting point for a dynamic theory of counterfactuals that is compatible with the dynamic approach here, and von Stechow (1999) discusses in some technical detail how counterfactuals depend on discourse-provided assumptions. I assume that any account that accurately describes presupposition projection in (42) will automatically generalize to the stupid examples in (39). That is, the failure of the presuppositions triggered by *stupid* to project in (39) is just an instance of the general behavior of any presupposition-triggering expression under the scope of counterfactual *would*.

4.7. Assessment

It would be hasty to conclude that the discussion in section 4 settles the question of whether stupid adjectives with infinitival complements have descriptive entailments. Undoubtedly there are alternative theories that get the truth conditions right for the cases discussed so far but which still allow for descriptive update for sentences like (43a):

- (43) a. It is rude to address the Queen.
b. It is rude to violate British royal protocol.

The idea would be that (43a) can conceivably be used descriptively to convey something generic about British royal protocol, in particular, that initiating a conversation with the Queen is against the rules. (43b), on the other hand, seems to be a much clearer case of pure sharpening, i.e., it is easier to decide that (43b) does nothing more than request the listener to make sure that their standard for rudeness is low enough that a certain type of behavior counts as rude.

What is at issue is whether the factors that determine the degree of rudeness of an event are presupposed to be part of the evaluation context. I claim that (43a) can inform the listener of descriptive facts only indirectly, by causing the listener to revise their beliefs about the nature of the presupposed event type (see the discussion of extenuating circumstances in section 4.3). If so, then (43a) can inform in the same way that a use of the sentence *It's a shame that Ann got pregnant* can be used to indirectly inform the listener that Ann is pregnant: through accommodation of a presupposition.

Of course it is the burden of anyone who wishes to argue that (43a) does have non-trivial descriptive entailments to go on to explain the unembeddability of stupid adjectives when they occur with an infinitival complements, as well as the similar unembeddability behavior of expressions such

surgically restored John's ability to feel physical pain: he would have stopped beating his dog as soon as it bit him.

as *did well to keep the hotel calm*.

In any case, even if relative uses of stupid adjectives do have some descriptive update potential under certain conditions, I take it that such uses are distinctly secondary to their typical uses as vagueness sharpeners. If so, then relative uses of stupid adjectives stand as a reasonably clear example of an expression type whose typical conventional use is primarily (if not exclusively) to bring about a change in the prevailing vague standards.

5. Conclusions

There are deep connections between traditional supervaluation treatments of vagueness and the dynamic perspective advocated here. In some sense, supervaluations approximate a dynamic semantics. Van Fraassen (1966) originally proposed supervaluation as a technique for dealing with truth-value gaps due to presupposition failure. Kamp, Fine, and others (see Fine (1975:fn. 13)) subsequently applied supervaluations to truth-value gaps due to vagueness: a sentence is super-true just in case it is true with respect to every possible way of resolving the vagueness of the predicates it contains. On the dynamic treatment, instead of considering ‘every possible way’ of resolving vagueness, only those ways of resolving vagueness provided by the discourse context are relevant. With contexts enriched with delineations, truth in the presence of vagueness is nothing more than entailment of a sentence by a context. Similarly, validity corresponds directly to standard dynamic validity as defined by, e.g., Groenendijk and Stokhof (1991:55). Thus one of the virtues of the dynamic approach is that it provides a unified dynamic account of truth and validity that simultaneously encompasses both presupposition and (one type of) vagueness.

So what does the dynamic approach have to say about the nature of vagueness? According to Pinkal (1983:47, commenting on Wittgenstein), “Knowing the meaning of an expression is knowing the ways the expression can be made more precise.” I have argued that it is the very act of using a vague expression that makes it more precise. Therefore we have arrived via a new route at the familiar conclusion that knowing what an expression means is knowing what a use of that expression does.

In order to develop an explicit theory of what a use of a vague expression does, I proposed that expressions containing gradable adjectives make a specific kind of contribution to contextual information. Accepting an utterance of, say, *Feynman is tall* can simultaneously refine our knowledge of Feynman’s height and of what counts as tall. An extension of the basic analysis predicts that update with a use of a measure phrase or a comparative is guaranteed to result in no sharpening of the relevant vague standard. In addition, the dynamic approach naturally leads to a reasonable account of higher-order vagueness in which modifiers such as *clearly* come out as vagueness quantifiers. A second extension of the basic analysis lead to the first detailed semantics for infinitival-taking adjectives. The somewhat surprising conclusion is that when certain adjectives such as *stupid* occur with an infinitival complement, the resulting construction has no update effect apart from its presuppositions and its effect on vague standards. Empirical support for this claim came from the otherwise unexplained phenomenon of unembeddability. In each case, the update effect of all of these expression types followed directly from stating their truth conditions in a dynamic framework.

References

- Ballweg, Joachim (1983) 'Vagueness or context-dependence? Supervaluation revisited in a semantics based on scales', in Thomas T. Ballmer and Manfred Pinkal (eds), *Approaching Vagueness*, North-Holland, Amsterdam, 59–78.
- Beaver, David (1997) 'Presupposition', in Johan van Benthem and Alice ter Meulen (eds), *Handbook of Logic and Language*, Elsevier, Amsterdam, 939–1008.
- Beaver, David (to appear) *Presupposition and Assertion in Dynamic Semantics*, CSLI Publications, Stanford.
- Bosch, Peter (1983) 'Vagueness is context-dependence', in Thomas T. Ballmer and Manfred Pinkal (eds), *Approaching Vagueness*, North-Holland, Amsterdam, 189–210.
- Bresnan, Joan (1972) *Theory of Complementation in English Syntax*, PhD Dissertation, MIT, Cambridge, MA.
- Cormack, Annabel (1998) *Definitions: Implications for Syntax, Semantics and the Language of Thought*, Garland, New York.
- Eikmeyer, Hans-Jürgen and Hannes Rieser (1983) 'A formal theory of context dependence and context change', in Thomas T. Ballmer and Manfred Pinkal (eds), *Approaching Vagueness*, North-Holland, Amsterdam, 131–188.
- Farkas, Donka F. (1988) 'On obligatory control', *Linguistics and Philosophy* **11**:27–58.
- van Fraassen, Bas C. (1966) 'Singular terms, truth-value gaps, and free logic', *Journal of Philosophy* **63.17**:481–95.
- Fine, Kit (1975) 'Vagueness, truth and logic', *Synthese* **30**:265–300.
- von Fintel, Kai (1999) 'Counterfactuals in a Dynamic Context', to appear in Michael Kenstowicz (ed), *Ken Hale: A Life in Language*, MIT Press.
- Gawron, Jean Mark (1995) 'Comparatives, superlatives, and resolution', *Linguistics and Philosophy* **18.4**:333–380.
- Groenendijk, Jeroen and Martin Stokhof (1990) 'Dynamic montague grammar', in *Papers from the Second Symposium on Logic and Language*, László and László Pólos (eds), Akadémiai Kiadó, Budapest, 3–48.
- Groenendijk, Jeroen and Martin Stokhof (1991) 'Dynamic predicate logic', *Linguistics and Philosophy* **14**:39–100.
- Hamann, Cornelia (1991) 'Adjectival semantics', in Arnim von Stechow and Dieter Wunderlich (eds), *Semantics: An International Handbook of Contemporary Research*, de Gruyter, New

York, 657–673.

Heim, Irene (1982) *The Semantics of Definite and Indefinite Noun Phrases*, PhD Dissertation, University of Massachusetts, Amherst, MA.

Heim, Irene (1983) ‘On the projection problem for presuppositions’, *WCCFL* 2, 114–125.

Heim, Irene (1985) ‘Notes on comparatives and related matters’, manuscript, UT, Austin, TX.

Heim, Irene (1992) ‘Presupposition projection and the semantics of attitude verbs’, *Journal of Semantics* 9:183–221.

Izvorski, Roumyana (1995) ‘A solution to the subcomparative paradox’, in *WCCFL XIV*, CSLI Publications, Stanford, 203–219.

Kamp, J. A. W. [Hans] (1975) ‘Two theories about adjectives’, in Edward Keenan (ed), *Formal Semantics of Natural Language*, Cambridge University Press, Cambridge, 123–155.

Kamp, Hans (1988) ‘Comments’, in *Contents of Thoughts*, R. Grim and D. Merrill, eds., University of Arizona Press, Tucson, Arizona, 156–181.

Kamp, Hans and Barbara Partee (1995) ‘Prototype theory and compositionality’, *Cognition* 57:129–191.

Karttunen, Lauri (1971) ‘Some observations on factivity’, *Papers in Linguistics*, 4.1:55–69.

Karttunen, Lauri (1974) ‘Presuppositions and linguistic context’, *Theoretical Linguistics* 1:181–194.

Kasper, Walter (1992) ‘Presuppositions, composition, and simple subjunctives’, *Journal of Semantics* 9:307–331.

Kennedy, Christopher (1997) *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*, PhD Dissertation, UCSC, Santa Cruz, CA.

Kennedy, Christopher (in press) ‘Cross-polar anomaly and the ontology of “degrees”’, *Linguistics and Philosophy*.

Kennedy, Christopher and Louise McNally (1999) ‘From event structure to scale structure: degree modification in deverbal adjectives’, in Mathews, T. and D. Strolovitch (eds), *SALT IX*, CLC Publications, Ithaca, 163–180.

Klein, Ewan (1980) ‘A semantics for positive and comparative adjectives’, *Linguistics and Philosophy* 4.1:1–45.

Klein, Ewan (1991) ‘Comparatives’, in Arnim von Stechow and Dieter Wunderlich (eds), *Semantics: An International Handbook of Contemporary Research*, de Gruyter, New York, 673–691.

- Kratzer, Angelika (1989) 'An investigation of the lumps of thought', *Linguistics and Philosophy* **12.5**:607–653.
- Kyburg, Alice and Michael Morreau (1996) 'Vague utterances and context change', AAAI Symposium, Technical Report FS-96-04, ISBN 1-57735-019-7.
- Kyburg, Alice and Michael Morreau (2000) 'Fitting words: vague language in context', *Linguistics and Philosophy* **23**:577–597.
- Lerner, Jan and Manfred Pinkal (1995) 'Comparative ellipsis and variable binding', in Mandy Simons and Teresa Galloway (eds), *Proceedings of SALT V*, Cornell University Linguistics, Cornell, 222-236.
- Lewis, David (1970) 'General semantics', *Synthese* **22**: 18–67; also in Davidson and Harmon (eds), *Semantics of Natural Language*, Reidel, Dordrecht, 169–218; also in his *Philosophical Papers*, vol. 1, 1983, Oxford University Press, Oxford.
- McConnell-Ginet, Sally (1973) *Comparative Constructions in English: A Syntactic and Semantic Analysis*, PhD Dissertation, University of Rochester, Rochester, NY.
- Moltmann, Frederike (1992) 'The empty element in comparatives', *NELS* 23, 319–333.
- Parikh, Rohit (1994) 'Vagueness and utility: the semantics of common nouns', *Linguistics and Philosophy* **17**:521–535.
- Pinkal, Manfred (1983) 'Towards a semantics of precization', in Thomas T. Ballmer and Manfred Pinkal (eds), *Approaching Vagueness*, North-Holland, Amsterdam, 13–57.
- Pinkal, Manfred (1989) 'Imprecise concepts and quantification', in R. Bartsch, J. van Benthem and P. van Emde Boas (eds), *Semantics and Contextual Expression*, Foris, Dordrecht, 221–265.
- Portner, Paul (1992) *Situation Theory and the Semantics of Propositional Expressions*, PhD Dissertation, UMass, Amherst, MA.
- Rivière, Claude (1983) 'Modal adjectives: transformations, synonymy, and complementation', *Lingua* **59**:1–45.
- Rullman, Hotze (1995) *Maximality in the semantics of WH-constructions*, PhD Dissertation, UMass, Amherst, MA.
- Sapir, Edward (1944) 'Grading, a study in semantics', *Philosophy of Science* **11.2**:93–116.
- Stalnaker, Robert (1973) 'Presuppositions', *Journal of Philosophical Logic* **2**:447–457.
- Stalnaker, Robert (1998) 'On the representation of context', *Journal of Logic, Language, and Information* **7**:3–19.

von Stechow, Arnim (1984) 'Comparing semantic theories of comparison', *Journal of Semantics* **3**:1–77.

Wilkinson, Robert (1970) 'Factive complements and action complements', *CLS* **6**: 425–44.

Wilkinson, Robert (1976) 'Modes of predication and implied adverbial complements', *Foundations of Language* **14.2**: 153–194.

Williamson, Timothy (1994) *Vagueness*, Routledge, London.

Williamson, Timothy (1999) 'On the structure of higher-order vagueness', *Mind* **108**:127–144.