

---

# Adjectival modification and gradation

Daniel Lassiter<sup>1</sup>

Department of Linguistics, Stanford University  
Margaret Jacks Hall, Building 460  
450 Serra Mall  
Stanford, California 94305 USA  
[danlassiter@stanford.edu](mailto:danlassiter@stanford.edu)

**Key words:** Adjectives, adverbs, modification, degree semantics, comparatives, superlatives, vagueness

---

A draft chapter for the Wiley-Blackwell *Handbook of Contemporary Semantics — second edition*, edited by Shalom Lappin and Chris Fox. This draft formatted on 21st May 2014.

## 1 Introduction

Adjectives modify nouns and noun phrases directly, and can be predicated of individuals, propositions, and events. This chapter deals with adjectival modification of nouns, and predicative constructions involving nouns and adjectives. Given the richness of the subject matter the treatment is necessarily selective; it is also, regrettably, confined mainly to English. However, it should provide enough background for further exploration of the literature and points of comparison for work on other languages. For a thorough empirical survey of English adjectives and a treatment of their interaction with adverbs, see Huddleston & Pullum 2002. For a theoretically-oriented survey covering the same material as this chapter in greater detail (and more), see Morzycki 2013.

Consider, for example, the adjective *beautiful*. In (1) it is used PREDICATIVELY, serving as the complement of the copular verb *to be*; the phrase *is beautiful* functions as the main predicate of the sentence.

(1) Prague is beautiful.

Looking just at (1), it's tempting to adopt a very simple treatment: there is some set of beautiful things in the world, and (1) says of Prague that it is an element of this set. Officially, then, the denotation of *beautiful* would be a semantic object of type  $\langle e, t \rangle$ , the characteristic function of the set of beautiful things. Assuming that the copula is semantically vacuous, the result is (2).<sup>1</sup>

(2)  $[[Prague\ is\ beautiful]]^{\mathcal{M}} = \mathbf{beautiful(Prague)}$

We might call this the ‘classical’ approach to adjective semantics, because it is in effect a generalization of the classical theory of concepts traditionally associated with Aristotle (see Murphy 2002: §2 for discussion).

A number of problems immediately arise with the classical approach. The first is that adjectives can also be used *attributively* as in (3), inside a noun phrase and, in this case, directly modifying a noun. If *beautiful* denotes a function from individuals to truth-values, we cannot account for this use without elaboration.

(3) Prague is a beautiful city.

---

<sup>1</sup> Assumptions and conventions: I use a semantics based on Montague 1973 as modified by Gallin (1975), where  $\mathcal{M} = \langle \mathbf{D}, [ \cdot ]^{\mathcal{M}} \rangle$  is a model consisting of a stratified domain of objects and an interpretation function which maps expressions into their model-theoretic interpretations. The basic types are (at least)  $s$  for worlds,  $e$  for individuals, and  $t$  for truth-values, and there is a function type  $\langle \alpha, \beta \rangle$  for any two types  $\alpha$  and  $\beta$ .  $\mathbf{D}$  is partitioned into subsets  $D_{\alpha}$  for each basic type  $\alpha$ . For any type  $\beta$ ,  $\mathbf{D}_{\beta}$  picks out the set of model-theoretic objects of type  $\beta$ .

I use *italicized* words and phrases to refer to natural language expressions, and **boldface** expressions to pick out their model-theoretic translations. I will mostly ignore intensionality, except where it is specifically relevant to a theoretical issue.

Assuming that *city* has the usual meaning of a common noun (type  $\langle e, t \rangle$ ), it is looking for an individual argument. Since *beautiful city* is a syntactic constituent in (3), we expect one to take the other as an argument; but if both sub-expressions are of type  $\langle e, t \rangle$ , neither can take the other as an argument and composition cannot proceed. §2 will discuss several approaches to this and related problems.

Second, some adjectives *can't* be used predicatively, and so the classical analysis would clearly be inappropriate for them. For example, the oddness of (4a) is presumably due to the fact that there is no set of 'former' things; yet (4b) is acceptable. This type of adjectival modification has been one source of inspiration for the intensional treatment of adjectives discussed in §2.

- (4) a. # Al is former. (predicative use unacceptable)  
 b. Al is a former politician. (attributive use OK)

A further puzzle is that *beautiful* and many other adjectives are GRADABLE: they do not appear only in the unmodified ('positive') form, but can form part of various morphologically and syntactically complex constructions which bear intricate logical relations to each other.

- (5) a. Prague is very/quite/sort of beautiful. (degree modifiers)  
 b. London is more beautiful than Manchester. (comparatives)  
 c. Of these cities, Barcelona is the most beautiful. (superlatives)  
 d. Paris is too beautiful for words. (excessives)

Gradability has important consequences for the way that we set up the semantics, even for simple cases like (1). To see why, consider that — if *beautiful* picks out a set — then *very beautiful* should pick out a smaller set (the particularly beautiful among the beautiful things). The denotation of the complex adjective phrase *beautiful but not very beautiful* should also be a proper subset of the denotation of *beautiful*. Obviously these two sets should be non-overlapping, and anything which falls into the first should count as *more beautiful* (5b) than anything in the second set. It is not at all clear how we could capture these logical relations among the positive, comparative, and various modified forms if the basic meaning of *beautiful* is an unordered set.

Along similar lines, consider (5b). Obviously, if this sentence is true then it cannot be that Manchester is beautiful while London is not. It could be, however, that both cities are beautiful, or that neither is. A theory of gradable adjectives should be able to capture these entailments across different constructions involving gradable adjectives (§§3-5).

A much-discussed feature of *beautiful*, *very beautiful*, and the like is their VAGUENESS and CONTEXT-DEPENDENCE. There is a vast literature on these issues. We will treat them fairly briefly, and in a way that emphasizes the compositional semantics of degree expressions, and their use in communication rather than broader logical or philosophical issues. See also Chapter 24 of this handbook for a cognitive science perspective on the semantics-pragmatics

interface which connects directly with the issues of vagueness and context-sensitivity discussed here.

§5 discusses scope interactions between degree expressions and quantifiers, modals, and other operators. In addition to providing a number of puzzles which are interesting in their own right, the open nature of this topic suggests that much work remains to be done in integrating degree semantics with other areas of natural language semantics, and that certain foundational assumptions may well need to be reconsidered in order to make progress on these difficult issues.

## 2 Adjective-noun combination

### 2.1 Kinds of adjectival modification

Clearly, a vegetarian farmer is someone who is both a vegetarian and a farmer, or (equivalently) someone who falls into the intersection of the set of vegetarians and the set of farmers. Adjectives like this are called INTERSECTIVE. By analogy, we might expect that a beautiful city is something that is both beautiful and a city; that a skilled craftsman is someone who is both skilled and a craftsman; and that a former friend is someone who is former and a friend. But these paraphrases become increasingly implausible as we move down the list, and the last is simply nonsensical. In fact these adjectives fall into various classes according to an influential typology of adjectives (Kamp, 1975; Siegel, 1976; Partee, 1995; Kamp & Partee, 1995).

Let  $A$  be an arbitrary adjective, and  $N$  an arbitrary noun. Intersective adjectives  $A$  are those for which the set of things that satisfy  $AN$  is simply the intersection of the set of things that satisfy  $A$ , and the set of things that satisfy  $N$ . For example, *vegetarian* is intersective because the vegetarian farmers are the vegetarians who are farmers, the vegetarian cellists are the vegetarians who are cellists, etc. Intersective adjectives thus license both of the patterns of inference in (6):

- (6) Al is a vegetarian farmer; Al is a cellist.  
 So, Al is a farmer.  
 So, Al is a vegetarian cellist.

SUBJECTIVE adjectives are a larger class which include the intersective adjectives. If  $A$  is a subjective adjective, then the set of things that satisfy  $AN$  is a subset of the things which satisfy  $N$ . For example, a skillful farmer is surely a farmer, but whether or not he counts as skillful depends on what kind of skill is under discussion — if it's farming skill, yes; if it's musical skill, we can't be sure until we learn more about him. As a result, being a skillful farmer and being a cellist is compatible with *not* being a skillful cellist. The signature feature of (non-intersective) subjective adjectives is thus the success of the first inference in (7) and the failure of the second.

- (7) Al is a skillful farmer; Al is a cellist.  
 So, Al is a farmer.  
 # So, Al is a skillful cellist.

Examples of NONSUBJECTIVE adjectives include *alleged*, *wannabe*, *fake*, and *former*. For these adjectives, the set of things which satisfy  $AN$  is *not* always a subset of the things which satisfy  $N$ . In general, then, if  $A$  is nonsubjective then there will be some things which satisfy  $AN$  without satisfying  $N$ : an alleged thief may or may not be a thief, and a wannabe actor probably isn't yet an actor. Inferences like (6) and (7) are not generally valid here either.

- (8) Al is an alleged forger; Al is a pickpocket.  
 # So, Al is a forger.  
 # So, Al is an alleged pickpocket.

A putative subclass of nonsubsective adjectives are the *privative* adjectives, which are marked out by the feature that something that satisfies *AN never* satisfies *N*. The classic examples are *fake*, *counterfeit*, *false*, and the like: the usual judgment is that a fake gun is not a gun, and a counterfeit dollar is not a dollar. However, Partee (1995, 2007) argues that there are no truly privative adjectives, pointing out (among other things) that the question ‘Is that gun real or fake?’ is not trivial as such an analysis would predict.

## 2.2 Intensional treatment

One prominent approach to these data is to adopt the Montagovian strategy of generalizing to the worst case. Instead of treating the simple predicative use of adjectives illustrated in (1) as basic, we begin with an account of the most complex cases we can find — such as *former* — and treat simple uses as special cases. More concretely, rather than treating adjectives as denoting the characteristic functions of sets of individuals, we will now analyze them as functions which take the intensions of nouns as arguments and return arbitrarily modified intensions (Montague, 1970; Kamp & Partee, 1995). For example, *alleged* would be analyzed as a function which maps noun intensions to derived noun intensions.

- (9) a.  $[[alleged]]^M = \lambda P_{\langle s, et \rangle} \lambda w_s \lambda x_e [\mathbf{alleged}(P)(w)(x)]$   
 b.  $[[alleged\ pickpocket]]^M = \lambda w_s \lambda x_e [\mathbf{alleged}(\mathbf{pickpocket})(w)(x)]$

The point of treating adjectives as functions on noun intensions is to block inferences like those in (8). Suppose that the meaning of *alleged* is a function which maps a noun meaning and a world to the set of individuals who have been said to be in the extension of the noun in that world. Clearly, nothing follows logically about whether the individual actually is in the extension, or about whether the individual has been said to be in the extension of any other noun. Another way to put the point is this: even if all and only forgers were pickpockets in our world, we wouldn’t be able to infer from *Al is an alleged forger* that Al is an alleged pickpocket. By operating on intensions, we can ensure that this inference fails, simply because there are alternative possible worlds in which *forger* and *pickpocket* do not have the same extension. By contrast, if *alleged* had an extensional and intersective meaning, this inference would be valid.

The intensional treatment makes room for complex meanings such as those expressed by *former* and *wannabe*, where the relationship between being *N* and *AN* cannot be expressed by simple mechanisms like set intersection. However, it also makes it necessary to introduce additional mechanisms to ensure that the valid inferences in (6)-(7) do go through, as well as other entailments.

The Montagovian strategy is to add meaning postulates which are lexically associated with the appropriate classes of adjectives. Some informal examples:

- (10) a. If  $A$  is in  $\{skilled, good, \dots\}$  then, for all  $x$ ,  $[[AN]]^{\mathcal{M}}(x)$  implies  $[[N]]^{\mathcal{M}}(x)$ .  
 b. If  $A$  is *former* then, for all  $x$ ,  $[[AN]]^{\mathcal{M}}(x)$  implies that there is some time  $t$  prior to the utterance time such that  $[[N]]^{\mathcal{M}}(x)$  holds at  $t$ .

Probably many more such rules would be needed, given the large and semantically varied adjective inventory of English.

This approach is pleasingly general, but it runs the risk of obscuring interesting details of the meanings of the adjectives in question. For example, Kamp & Partee (1995) point out that some apparently non-intersective adjectives may be better explained as having context-dependent but intersective meanings. It would be unwise to draw the inference in (11), for example; but this is probably not due to *tall* being non-intersective, but rather a subtle shift in meaning induced by combining the adjectives with different nouns.

- (11) Al is a tall jockey; Al is a hockey player. # So, Al is a tall hockey player.

This, in turn, is presumably related to facts about the different distributions of heights among jockeys and hockey players, a piece of contextual information which influences the interpretation of adjectives like *tall*. Note in favor of this analysis that the inference is reasonable if we reverse the nouns. Intuitively, this is because hockey players tend to be taller than jockeys, so that someone who is tall *for a hockey player* is probably also tall *for a jockey*.

- (12) Al is a tall hockey player; Al is a jockey. So, Al is a tall jockey.

If this is right, it may be possible to analyze *tall* as an intersective adjective after all. The methodological lesson is that we must be careful to hold the context fixed in applying tests such as (11).

Along similar lines, it may be possible to treat the failure of (7) with the classic non-intersective adjective *skillful* as being due to the presence of an implicit argument specifying which kind of skill is relevant to its interpretation. On this treatment, the failure of (7) would not show that *skillful* is non-intersective, but that subtle shifts in its meaning are induced by the change in the context between the first premise and the conclusion: in other words, (7) is intuitively invalid because it is interpreted like (13).

- (13) Al is skillful *as a farmer*; Al is a cellist. # So, Al is skillful *as a cellist*.

Making this precise would, of course, require an account of explicit *as*-phrases as well as their putative implicit counterparts.

Another difficulty for the intensional treatment is that it is not obvious how to account for predicative uses of adjectives. What is needed is an explanation of the fact that intersective and subsective adjectives can usually be used predicatively, but nonsubsective adjectives frequently cannot.

- (14) a. Al is vegetarian. (intersective)  
 b. Al is skillful. (subsective)  
 c. # Al is alleged. (nonsubsective)

In an important early treatment, Siegel (1976) argues that the attributive and predicative uses of adjectives like *vegetarian* and *skillful* are really different lexical items, even if they happen to be homophonous in many cases. On this account the problem with (14c) comes down to a lexical gap in English. Partee (1995) suggests rather more parsimoniously that intersective adjectives are listed in the lexicon as simple predicates and operated upon by a general type-shifting rule in their attributive uses. (15) gives a simple-minded implementation of this idea.

- (15) a.  $[|vegetarian|]^{\mathcal{M}} = \lambda w_s \lambda x_e [\mathbf{vegetarian}(w)(x)]$  (basic meaning)  
 b.  $\mathbf{ATT} = \lambda P_{\langle s,et \rangle} \lambda Q_{\langle s,et \rangle} \lambda w_s \lambda x_e [P(w)(x) \wedge Q(w)(x)]$  (type-shifter)  
 c.  $[|vegetarian\ farmer|]^{\mathcal{M}} = \mathbf{ATT}([|vegetarian|]^{\mathcal{M}})([|farmer|]^{\mathcal{M}})$   
 $= \lambda w_s \lambda x_e [\mathbf{vegetarian}(w)(x) \wedge \mathbf{farmer}(w)(x)]$

We could then explain the unacceptability of *Al is alleged/former* by treating these adjectives as being listed in the lexicon in the higher (property-modifying) type, so that they cannot apply directly to an individual.

If adjectives like *skillful* are really intersective but context-dependent, as we speculated above, then this treatment may be able to account for the data we have seen so far. However, there must still be room for lexical restrictions on the availability of attributive and predicative uses of adjectives, even intersective ones: compare *The baby is asleep* to the rather less natural (but attested) *the asleep baby*.

### 2.3 Modification of individuals and events

Above we suggested that *old* and *skillful* might be intersective after all, once certain non-obvious features of their meaning are taken into account. Larson (1998) argues in a somewhat different way that some or all apparently intensional/nonintersective adjectives can be treated as extensional and intersective. He focuses in particular on the fact that many adjectives are ambiguous between an intersective and a non-intersective reading, as in the famous example *Olga is a beautiful dancer*. This sentence has two very different readings:

- (16) a. ‘Olga is beautiful, and she is a dancer.’ (‘intersective’)  
 b. ‘Olga dances beautifully.’ (‘non-intersective’)

On reading (16a) the sentence entails that Olga is beautiful, but leaves open that her dancing could be atrocious; on reading (16b) it entails that her dancing is beautiful but does not exclude the possibility that she is quite ugly as a person. Many other adjectives display similar ambiguities, including *skillful*, the showcase subsective adjective in the previous section: *Al is a skillful*

*farmer* is most naturally interpreted as meaning that Al is skillful *as a farmer*, but could also be used to mean that he is both a farmer and skillful at some other salient activity.

Larson points out that the substitution failures that the intensional treatment is designed to account for also occur in similar constructions for which intensionality is not an obvious diagnosis. Note first that, on the non-intersective reading, the inference in (17) fails.

(17) Olga is a beautiful dancer; Olga is a singer. # So, Olga is a beautiful singer.

This is quite similar to:

(18) Olga dances beautifully; Olga sings. # So, Olga sings beautifully.

A reasonable diagnosis of the substitution failure in (18) is that manner adverbs like *beautifully* are modifiers of EVENTS. A standard event semantics (Davidson, 1967; Parsons, 1990) predicts the failure of (18) in simple first-order terms. Very roughly (letting  $e$  be a variable over events, and glossing over important but mostly orthogonal issues about the choice of quantifier):

(19) a.  $[[\textit{Olga dances beautifully}]^{\mathcal{M}}] = \exists e[\mathbf{dancing}(e, \mathbf{Olga}) \wedge \mathbf{beautiful}(e)]$

b.  $[[\textit{Olga sings}]^{\mathcal{M}}] = \exists e[\mathbf{singing}(e, \mathbf{Olga})]$

c.  $[[\textit{Olga sings beautifully}]^{\mathcal{M}}] = \exists e[\mathbf{singing}(e, \mathbf{Olga}) \wedge \mathbf{beautiful}(e)]$

Clearly, (19a) and (19b) can be true while (19c) is false. Larson points out that the non-intersective reading of *Olga is a beautiful dancer* can be treated similarly, assuming plausibly that the meaning of the deverbal noun *dancer* contains an event variable which can be modified by the adjective. The two readings of this sentence are then generated by allowing the adjective *beautiful* to modify either an individual variable or an event variable. The two readings of *beautiful dancer* come out as in (20).

(20) a.  $\lambda x \exists e[\mathbf{dancing}(e, x) \wedge \mathbf{beautiful}(x)]$  ('intersective')

b.  $\lambda x \exists e[\mathbf{dancing}(e, x) \wedge \mathbf{beautiful}(e)]$  ('non-intersective')

Note that *beautiful* is a simple predicate in both cases, and differs only in what it is predicated of; Larson's point is that we can treat some apparently non-intersective adjectives as having simple intersective meanings if we make sure that we are correctly representing their interactions with the meaning components provided by the the noun.

This line of attack may even succeed with some nonsubsective adjectives. For example, Larson suggests treating *Al is a former teacher* as having two readings as well.

(21) a.  $\exists e[\mathbf{teaching}(e, \mathbf{Al}) \wedge \mathbf{former}(\mathbf{Al})]$  ('intersective')

b.  $\exists e[\mathbf{teaching}(e, \mathbf{Al}) \wedge \mathbf{former}(e)]$  ('non-intersective')

Potentially, the individual-modifying ('intersective') reading is not available simply because it makes no sense: unlike *beautiful*, *former* picks out a property that can only be predicated sensibly of events. This would account for the fact that only reading (21b) is available, and also for the fact that *former* cannot be used predicatively (cf. (14)).

It remains to be seen whether the full range of adjectives can be treated in this way. Larson points out that some attributive-only adjectives such as *mere* and *utter* are not plausibly treated as event modifiers, and suggests that the nouns that these adjectives combine with may have still further intricacies. Non-subsective adjectives such as *alleged* and *fake* may also pose challenges to the approach. Overall, Larson's approach suggests that it may be possible to simplify the typology of adjectives while also drawing connections with the semantics of nominals and the syntax of DPs. However, a good deal of empirical and theoretical work remains to be done in order to make good on this promise, requiring simultaneous consideration of evidence and theoretical issues from morphology, syntax, semantics, and pragmatics as well as the issues involving gradation and scales discussed in the remainder of this chapter.

### 3 Gradation and scales

#### 3.1 Diagnosing gradability

Many of the adjectives that we have discussed are GRADABLE, including *beautiful*, *old*, *tall* and *skilled*. Diagnostics for gradability include participation in comparative and equative constructions; the availability of complex constructions involving degree modifiers and measure phrases; the possibility of using overt comparison classes. For example:

- (22) a. This car is older than that one. (comparatives)  
 b. This car is as old as that one. (equatives)  
 c. This car is very/quite/somewhat old. (degree modification)  
 d. This car is ten years old. (measure phrases)  
 e. This car is old for a Honda. (comparison classes)

A particular gradable expression may not participate in all of these constructions for principled or idiosyncratic reasons. For example, *beautiful* differs from *old* in not readily accepting measure phrases; presumably, this has to do with the fact that beauty is not easily measured or associated with conventional units. Extreme adjectives such as *outstanding* and *enormous* also show complicated patterns that are not well understood: consider for example *quite/?very/#somewhat enormous* and the contrast in (23).

- (23) a. # That house is more enormous than this one.  
 b. This house is enormous, and that one is even more enormous.

Some of the adjectives we have considered — *vegetarian*, *alleged*, *wannabe*, and *former* — fail all of the tests in (22) and are apparently non-gradable. This is interesting because each of these adjectives is associated with a kind of meaning which *could* sensibly be graded. For example, from the nature of the concepts involved we might expect them to form comparatives with the meanings paraphrased in (24).

- (24) a. Sam is more vegetarian than Al. ('Al eats more meat')  
 b. Bill is a more alleged thief than Mary. ('More people say this of Bill')  
 c. Al is a very former teacher. ('He stopped teaching a long time ago')

The fact that the sentences in (24) are quite odd suggests that grammatical gradability may not be straightforwardly predictable from the nature of the property that an adjective expresses. It may be necessary instead to simply supply certain adjectives with an extra argument (degree, comparison class, or both) which can be manipulated by operators such as degree modifiers and comparatives. Non-gradable adjectives, then, would simply be adjectives which lack this additional argument and are listed as properties or functions over properties (as in §2). This treatment of the gradable/non-gradable distinction is less than fully satisfying — it would be much nicer to have a uniform semantic treatment of adjectives in which non-gradable adjectives emerge as a

special case without lexical stipulation — but data such as (24) provide a certain amount of empirical motivation for a lexical distinction. We will assume it in what follows.

Note that gradability is not limited to the syntactic category of adjectives. In English, gradable expressions occur in various syntactic categories including adverbs (*very/more quickly*), verbs (*love Sam more than Mary does*), quantificational determiners (*more water/boys*), nouns (*be more of an artist than Bill*) and auxiliary modals (*You should leave more than he should*). See e.g. Solt 2009; Wellwood *et al.* 2012; Lassiter 2014 for data, theoretical discussion, and further references on gradation beyond adjectives.

The existence of far-reaching gradation in natural languages really should not be a great surprise: most psychologists who study concepts have long since abandoned the classical assumption that concepts have sharp boundaries in favor of graded representations of concepts, expressed using (e.g.) probability theory, fuzzy logic, or vector spaces. See Murphy 2002 for a history of the transition from classical to graded theories of concepts and a survey of relevant theory and experimental data, including his §2 on empirical problems for the classical theory and §11 on the relationship between concepts and word meanings.

### 3.2 Modeling gradability with and without degrees

Historically there have been two main approaches to the semantics of gradation. Bartsch & Vennemann 1973 and many following theories proceed by adding an extra semantic type  $d$  for DEGREES. Degrees are abstract representations of measurement (such as heights and weights) and come organized into ordered sets called SCALES. Formally, a scale is a structure  $\langle D, \geq \rangle$ , where  $D$  is a set of degrees and  $\geq$  is a reflexive, transitive, antisymmetric, and possibly connected binary order on  $D$ . (I presuppose here some basic concepts from order theory, at the level of Partee *et al.* 1990, §3.) As we will see below, some scales seem to have richer logical structure; but this is the minimum.

On this style of analysis, gradable expressions are provided with a degree argument which can be bound by operators such as comparatives and measure phrases. Non-gradable and gradable adjectives thus differ in their semantic type: the former are (mostly) simple predicates (type  $\langle e, t \rangle$ ), while the latter are functions from a degree to a predicate (type  $\langle d, et \rangle$ ).<sup>2</sup> For example, the lexical entry for *tall* in such a theory would be as in (25a), and the measure

<sup>2</sup> Some clarificatory notes and pointers to the literature: [1] There is a debate about whether the right semantic type for gradable adjectives is  $\langle e, d \rangle$  or  $\langle d, et \rangle$ ; see Bartsch & Vennemann 1973; Kennedy 1997 and Cresswell 1976; von Stechow 1984; Heim 2001 respectively. Here we will focus on the latter. [2] An additional type may be needed for intensional adjectives unless the strategy of Larson (1998) to reduce them to intersective adjectives is successful (§2.3). [3] Modal adjectives such as *likely* take propositional arguments and thus will be of type  $\langle d, \langle \langle s, t \rangle, t \rangle \rangle$  rather than  $\langle d, et \rangle$ . Control adjectives such as *eager* presumably have an even more

phrase *five feet* picks out a measure of height which saturates the degree variable in the meaning of *tall*, returning a property which is true of anything which is at least 5 feet tall.

- (25) a.  $[[tall]]^{\mathcal{M}} = \lambda d_d \lambda x_e [\mathbf{height}(x) \geq d]$   
 b.  $[[five\ feet]]^{\mathcal{M}} = 5'$   
 c.  $[[five\ feet\ tall]]^{\mathcal{M}} = \lambda x_e [\mathbf{height}(x) \geq 5']$

The main alternative is to suppose that gradable and non-gradable adjectives share the semantic type  $\langle e, t \rangle$ , but differ in that only gradable adjectives are semantically context-sensitive. On one version of this approach (Lewis, 1970; Barker, 2002), comparatives, measure phrases, and other degree operators are treated not as expressions which bind a degree variable, but rather as expressions which shift a contextual degree (or ‘delineation’) parameter that controls the interpretation of a gradable adjective. For example, instead of (25) we would have (26), with the ‘shifty’ entry (26c) for the measure phrase.

- (26) a.  $[[tall]]^{\mathcal{M}, \Delta} = \lambda x_e [\mathbf{height}(x) \geq d_{tall}]$   
 b.  $[[five\ feet]]^{\mathcal{M}, \Delta} ([[A]]^{\mathcal{M}, \Delta}) = ([[A]]^{\mathcal{M}, \Delta [d_A \leftarrow 5']})$   
 c.  $[[five\ feet\ tall]]^{\mathcal{M}, \Delta} = \lambda x_e [\mathbf{height}(x) \geq 5']$

In these definitions,  $\Delta = \langle d_{tall}, d_{happy}, \dots \rangle$  is a very long list of delineations, one for each scalar adjective in the language.  $\Delta [d_A \leftarrow 5']$  is the list which is everywhere identical to  $\Delta$  except that  $d_A$  is replaced by  $5'$ .

The difference between (25) and (26) is essentially whether there is object-language quantification over degrees. In fact, a parametric semantics like (26) can always be rewritten using object-language quantification over the parameters as in (25) (cf. Cresswell, 1990). Although the proposals in (25) and (26) are superficially different, they are really semantically equivalent. (They are not equivalent in their morphosyntactic predictions, though; see §3.3.)

Context-sensitive predicate analyses of gradable adjectives that differ more deeply from degree-based treatments have been offered by Klein (1980); Doetjes *et al.* (2009); Burnett (2012) and others. These analyses treat gradable adjective meanings as being relativized not to a degree parameter but to a comparison class parameter. In such approaches it is necessary to impose strong restrictions on possible context-sensitive meanings. For instance, if there is a possible context in which Al counts as ‘tall’ and Bill does not, there should not be any possible context in which Bill counts as ‘tall’ and Al does not. After all, if the former is true then Al is *taller* than Bill.

---

complex type. The denotations given below for degree operators will be specific to adjectives which take individual arguments, but are readily generalized to modal adjectives (Lassiter, 2014, §1). [4] A number of theorists have proposed treating degrees not as points on scales but as intervals (e.g. Kennedy, 2001; Schwarzschild & Wilkinson, 2002). The issues involved are fairly intricate and go beyond what we can consider in this space.

Putting together the various qualitative restrictions that are needed in order to avoid such monstrosities, it turns out that the meaning of any adjective *Adj* relies on a reflexive, transitive, possibly connected binary order (weak order)  $\succeq_A$  ‘at least as *Adj* as’ defined over a domain of individuals  $D_A \subseteq D_e$  (Klein, 1980). If the order is connected, the structure  $\langle D_A, \succeq_A \rangle$  is an ORDINAL SCALE, a type of qualitative scale which has received much attention in Measurement Theory (van Benthem 1983; Klein 1991; cf. Krantz *et al.* 1971). Standard techniques from Measurement Theory reveal that for any semantics of this form there is an equivalent degree-based semantics as in (25) built on a scale  $\langle D, \succeq \rangle$  (indeed, an infinite number of them). Similar correspondences hold for a number of more restrictive qualitative structures which are plausibly relevant to gradation in natural languages (Sassoon, 2010a; van Rooij, 2011a; Lassiter, 2014). Even explicit reference to degrees can be analyzed in qualitative terms by treating degrees as equivalence classes of individuals under the  $\succeq_A$  relation (Cresswell, 1976; Rullmann, 1995; van Rooij, 2011a, etc.).

Despite the rather different overt appearance, then, there does not seem to be much to choose in this debate either: insofar as they are rich enough to capture the basic phenomena, semantic treatments of gradable adjectives can be written equivalently using a degree-based semantics or a qualitative semantics. In giving a semantics for English adjectives, the choice between the two approaches is largely one of taste and ease of use.

### 3.3 Morphosemantics of the positive form

Even though the correspondence between theories with and without degrees is closer than it appears on the surface, there are good arguments in favor of theories treating the positive form as a predicate, coming not from semantics but from morphology. In degree-free theories a gradable adjective such as *tall* takes one argument, an individual. This means that sentences with the positive form of the adjective can be treated in a maximally simple way:

$$(27) \quad [Al \text{ is tall}]^{\mathcal{M}, \Delta} = \mathbf{height}(\mathbf{Al}) \geq d_{tall}$$

The effect is that *Al is tall* means ‘Al is at least as tall as some contextually specified height’.

In contrast, if we attempt to apply the denotation of *tall* in (25a) directly to an individual we get a type-mismatch, since the adjective expects a degree as its first argument. Degree-based theories generally deal with this problem by assuming that there is a phonetically null degree morpheme called “*pos*”. A gradable adjective must combine with *pos* before it can take an individual argument (Cresswell, 1976; von Stechow, 1984; Kennedy, 1997, etc.).

$$(28) \quad \begin{array}{l} \text{a. } [pos]^{\mathcal{M}, \Delta} = \lambda A_{\langle d, et \rangle} \lambda x_e [A(d_A)(x)] \\ \text{b. } [pos \text{ tall}]^{\mathcal{M}, \Delta} = \lambda x_e [\mathbf{height}(x) \geq d_{tall}] \\ \text{c. } [Al \text{ is pos tall}]^{\mathcal{M}, \Delta} = \mathbf{height}(\mathbf{Al}) \geq d_{tall} \end{array}$$

The “standard degree”  $d_{tall}$  in (28) performs the same function as the delineation parameter  $d_{tall}$  in (27); the difference is in whether the degree variable in the adjective’s meaning is automatically bound to this contextual parameter, or variable and parameter are connected by the action of *pos*.

As Klein (1980) points out, on this theory it is essentially an accident of English that the *pos* morpheme is silent. To the extent that gradable adjectives in other languages are also of type  $\langle d, et \rangle$ , we should expect to find overt counterparts of *pos* doing the same job in other languages. There are no clear candidates (though see Liu 2010 for an argument that Mandarin has such a morpheme, and Grano 2012 for a rebuttal). More recently Bobaljik (2012) argues using patterns of syncretism from a large sample of languages that the comparative form of adjectives universally contains the positive form. This is also potentially troubling for the *pos*-based theories, since they treat the comparative as containing not the full complex  $[pos\ A]$  but only the adjectival root  $A$ . However, the argument depends heavily on theoretical assumptions about morphological syncretism associated with Distributed Morphology (Halle & Marantz, 1993), and a defender of *pos* could perhaps appeal to a different theory of morphology in order to explain why Bobaljik’s patterns would emerge.

An undesirable feature that is shared by the *pos*- and delineation-based theories is that they require the interpretation of English sentences to be relativized to a huge number of parameters — one delineation or standard degree for each adjective in the language, whether or not the relevant adjective actually appears in the sentence. One possible way to avoid this, while also avoiding the morphological problems of the *pos*-based theory, is to suppose that *pos* is not a morpheme but an instantiation of type-shifting mechanisms which are freely available and quite generally phonologically unrealized. Briefly, the idea would be that sentences do not necessarily denote propositions, but may denote functions from a small set of arguments to propositions; interpreters then use pragmatic and world knowledge to fill in appropriate values for the unsaturated variables. *Sam saw her*, for example, would denote a function  $\lambda x.\text{saw}(x)(\mathbf{Sam})$ , and an interpreter must fill in a value for the unsaturated variable in order to recover a proposition. Such a theory must rely heavily on type-shifting mechanisms which intervene to allow composition to proceed when it would otherwise halt. It must also supply a pragmatic story about how interpreters infer appropriate values of unsaturated variables, and when speakers can reasonably assume that listeners will be able to perform this task (see Goodman & Lassiter (this volume) for relevant discussion).

For the interpretation of positive-form adjectives, what would be needed in the simplest case is a type-shifter **POS** which simply reverses the order of the arguments of *tall*. Note that we no longer need to relativize interpretation to a long list of dedicated contextual parameters, one for each adjective of the language.

- (29) a.  $[[tall]]^M = \lambda d_d \lambda x_e [\text{height}(x) \geq d]$   
 b. **POS** =  $\lambda A_{\langle d, et \rangle} \lambda x_e \lambda d_d [A(d)(x)]$

$$c. \llbracket Al \text{ is tall} \rrbracket^{\mathcal{M}} = \mathbf{POS}(\llbracket tall \rrbracket^{\mathcal{M}})(\llbracket Al \rrbracket^{\mathcal{M}}) = \lambda d_d[\mathbf{height}(Al) \geq d]$$

Listeners must infer a reasonable value for the unsaturated degree variable in order to interpret (29c), just as they would for the contextual parameters in order to interpret (27) and (28c).

Whether this modification represents a genuine explanation of the silence of *pos*/**POS** or a mere terminological shift depends on the details of the compositional semantic theory in which it is embedded. In the context of a theory in which such type-shifting mechanisms are strongly motivated (Szabolcsi, 1987; Steedman, 1987, 2001; Jacobson, 1999; Barker & Shan, 2014), this approach may represent a genuine theoretical explanation which enables us to maintain other desirable features of the degree-based theory — in particular, a simple treatment of comparatives and their interactions with modals and quantifiers as discussed below.

Given the difficulty of finding clear empirical differences between the various ways of setting up the semantics of gradable adjectives, the most efficient route at this point seems to be to simply choose one and work with it. The rest of our discussion will assume an explicit degree-based semantics in the tradition of Bartsch & Vennemann (1973); von Stechow (1984); Bierwisch (1989); Kennedy (1997); Kennedy & McNally (2005). This choice is motivated chiefly by the fact that the degree-based semantics is somewhat simpler to state and work with, and it is important to keep in mind that there are many alternative ways to set up the semantics of these constructions, with and without degrees, which frequently generate very subtle empirical differences or none at all.

Note, however, that this approach is largely motivated by the current focus on adjectives in English. Beck *et al.* (2009); Bochnak (2013) argue that languages may vary in whether they make use of degrees in their semantics. On Bochnak's account, a degree-based semantics is needed for English; but the comparison-class-based semantics of Klein (1980) is essentially correct for the native American language Washo, which has vague scalar expressions but no direct comparatives or other degree-binding operators. This opens up the possibility that the existence of degrees could be motivated in languages with rich degree morphology by an indirect argument, by appealing to typological distinctions among languages that are difficult to explain if all languages represent and use degrees in the same way. Since the argument invokes a parametric distinction, it predicts a sharp discontinuity between languages with and without traces of degree binding. An alternative hypothesis is that languages may display gradual variation in the number of degree-binding operators in their lexicon, with English and Washo merely representing extremes of “many” and “zero” such items, respectively. More work is needed to clarify the empirical situation here.

### 3.4 Vagueness and context-dependence of the positive form

In §3.3 we discussed details of the compositional semantics and morphology of the positive form of gradable adjectives. These adjectives have also been the

primary subject matter of a vast literature on VAGUENESS, with important contributions from philosophers, linguists, psychologists, and computer scientists. We cannot hope to cover the intricate debates on vagueness in detail, but will settle instead for a quick overview of empirical characteristics and theories which interface with the compositional semantics described above. Some good entry points into the larger debates are Williamson 1994; Kamp & Partee 1995; Keefe & Smith 1997; Keefe 2000; Barker 2002; Shapiro 2006; Kennedy 2007; van Deemter 2010; van Rooij 2011b.

### Empirical feature of vague adjectives

The most fundamental diagnostic for vagueness is a LACK OF SHARP BOUNDARIES. That is, to say that *heavy* is vague is to say that we cannot identify a weight  $w$  such that anything that weighs  $w$  kilograms or more is heavy and anything that weighs less is not heavy. This is true even though there are clear cases of heavy things (a truck) and clear cases of things that are not heavy (a feather). A closely related characteristic is TOLERANCE (Wright, 1976). Suppose we have identified something that is definitely heavy. To say that *heavy* is tolerant is to say that we should not also identify something that is just a tiny bit lighter (something 1 microgram lighter than the truck) as not being heavy.

Unfortunately, acquiescing to these claims about *heavy* leads straight to the SORITES PARADOX.

- (30) a. This truck is heavy.  
 b. It's not the case that an object 1 microgram lighter than a heavy object is not heavy.  
 c. This feather is heavy.

If we can find an object exactly 1 microgram lighter than the truck, (30b) requires that it can't count as 'not heavy'; if we're working with classical logic, this means that the lighter object is heavy. We then find another object 1 microgram lighter than that, apply (30b) again and conclude that it is also heavy. Continuing this procedure for some tedious length of time, we will eventually reach a weight which is less than or equal to the feather's weight, from which we can conclude that the feather is heavy as well. So, if the truck is heavy and the tolerance principle is true of *heavy*, then the feather is heavy; but that is obviously false.

Vague adjectives also admit of BORDERLINE CASES. Let's allow that the truck is heavy and the feather is not; what about this table? Indeed, for virtually any context and purpose we can imagine there will be items for which it is unclear whether or not they count as 'heavy'. Note that there are (at least in principle) two ways that this could be spelled out. First, a borderline case of 'heavy' could be an object such that it is partly, but not fully, acceptable to describe it as 'heavy', and acceptable to a similar degree to describe it as 'not heavy'. Second, a borderline case could be an object

which falls into the range on a scale in between two regions of clarity, but for which both descriptions are clearly *inappropriate*. Whether both of these possibilities are instantiated is a matter for empirical investigation; see e.g. Égré (2011).

Vague adjectives typically display considerable CONTEXT-DEPENDENCE. This comes in at least two forms. STATISTICAL context-dependence involves the fact that the way that a property is distributed among other relevant objects can influence the truth-value of a sentence containing a relative adjective. For example, a house listed for \$400,000 might well count as expensive in Atlanta, where the average sale price of homes (at the time of writing) is less than \$200,000. A house with the same price would probably not count as expensive if it were being sold in San Francisco, where homes typically sell for around \$450,000. What varies between the two cases is the statistical distribution of prices among homes in the local area.

This kind of implicit relativization to a class of relevant comparisons is related to the linguistic phenomenon of overt comparison classes (Solt, 2011). The sentences in (31) might well be true of the same piece of property. Presumably this is because (31a) explicitly invokes the distribution of prices among condos, without excluding from consideration condos in cheaper locations; while (31b) invokes the distribution of prices among homes of all kinds in San Francisco, which is generally rather higher than the prices of condos nationwide.

- (31) a. This property is expensive for a condo.  
 b. This property is not expensive for a home in San Francisco.

A related but possibly different source of context-dependence relates to the goals and interests of the conversational participants (or other relevant people). For example, whether I consider a home with the prices quoted above to be expensive might depend not only on the objective statistics of a relevant comparison class, but also on what I can afford; a \$500,000 home may appear as expensive to someone who is very poor, and as not expensive to someone who is very rich, regardless of the local statistics of prices. This kind of context-sensitivity is subtler and less well-understood, but see Fara 2000 for an insightful discussion. Perhaps statistical context-sensitivity can even be reduced to interest-relativity, if it can be shown that statistical facts matter only when they pertain to the practical interests of the parties in a conversation.

## A sampling of theories

There are many theories of vagueness; here I will describe informally a small number which interface clearly with the degree-based semantics for gradable adjectives described above. I also won't stop to explain how each theory deals with the critical sorites paradox; see the works cited for extensive discussion.

The simplest approach to vagueness, in a certain sense, is to deny that it has anything to do with meaning. That is, we interpret *heavy* as a property

of objects whose weight exceeds some standard degree/delineation parameter  $d_{heavy}$ , and assume that this value is given with precision by the interpretation of the language together with facts about the context of use. What generates vagueness, on this view, is some kind of irresolvable uncertainty about what the correct interpretation of the language is. This is the EPISTEMIC THEORY of vagueness, defended most prominently by Williamson (1994). It has the undeniable advantage of keeping the semantics simple, and Williamson also argues that the theory is plausible given general epistemic limitations of humans. However, many linguists would balk at the rather extreme form of semantic externalism that Williamson's theory presupposes. One influential position would take it as absurd to suppose that there are facts about a language that speakers of the language do not and cannot know (Chomsky, 1986).

A related but perhaps less contentious idea relies on the fact that vague expressions are context-dependent, and that conversational participants may be uncertain about the precise nature of the conversational context — say, whether they are speaking in context  $c$  or context  $c'$  (or in context  $c''$  or ...). On this account, speakers know all of the relevant facts about their language, including what the precise linguistic interpretation *would be* if they were in  $c$  (or in  $c'$  or ...). Vagueness can then be modeled as uncertainty, as in the epistemic theory, but without requiring that there be a precise but unknown *linguistic* fact about how heavy an object must be in order to count as 'heavy'. Instead, the language provides a linguistic 'hook' for the context to fill in a value, but says nothing about what the value is. (Such an account does, however, imply that there would be no vagueness if all relevant facts about the context were fully known. One might reasonably doubt the plausibility of this consequence.)

One example of an account along these lines is Barker 2002. Barker points out that treating vagueness as uncertainty about the context allows us to capture many of the useful features of SUPERVALUATIONAL theories of vagueness (Fine, 1975; Kamp, 1975; Keefe, 2000) without building in special linguistic devices for managing uncertainty about the denotations of vague expressions. On the supervaluational account, the interpretation function (relative to a context) associates vague predicates with a range of 'admissible precisifications' for a vague adjective such as *tall*. For example, it might be that all and only values between 10 and 20 kilograms are admissible precisifications for 'heavy'. We can then say that ' $x$  is heavy' is clearly true if  $x$  weighs more than 20 kilos, since the sentence comes out as true under all admissible precisifications; and that it is clearly false if  $x$  weighs less than 10 kilos, since it comes out as false under all admissible precisifications. If  $x$ 's weight is between 10 and 20 kilos, then  $x$  is a borderline case of 'heavy' — that is, it counts as 'heavy' under some but not all admissible precisifications. Theories of this sort are able to account for the tolerance of predicates like *heavy* by extending classical logic. As long as the range of admissible precisifications is not too small, we will never move from 'heavy' to 'not heavy' in a single small

step. We will sometimes move from ‘heavy’ to ‘borderline’ in a single small step, but this is not a violation of the tolerance principle as it was formulated above.

Barker points out that a similar effect can be achieved by supposing that the interpretation function assigns a unique interpretation to *heavy* relative to any given context, but that there is uncertainty about what the relevant context is. That is, there is some set of epistemically possible contexts  $C$ , and the conversation might for all we know be taking place in any  $c \in C$ , each of which determines a delineation  $d_{heavy}$ . Suppose that we know enough about the context to exclude delineations greater than 20 kilos or less than 10 kilos. Then, by plugging in the delineation semantics discussed in §3.2, we have the result that ‘ $x$  is heavy’ is definitely true if  $x$ ’s weight is greater than 20 kilos, in the sense that we know that it will be true no matter what further information we acquire about the context of conversation. Similarly, ‘ $x$  is heavy’ will be definitely false if  $x$ ’s weight is less than 10 kilos. However, if  $x$ ’s weight falls between 10 and 20 kilos, we will be uncertain about whether  $x$  counts as ‘heavy’ since the context could (for all we know) determine a value for  $d_{heavy}$  which is greater or less than  $x$ ’s weight. Barker’s version of the theory is also able to account for the tolerance of *heavy*, if the principle is re-formulated to make reference to what is known about the context-sensitive meaning of *heavy* — that is, what is true no matter which  $c \in C$  turns out to be the true context. The revised principle requires that there be no  $w$  and small  $\epsilon$  such that something which weighs  $w$  kilos is known to be heavy and something which weighs  $w - \epsilon$  is known to be not heavy.

Both classical supervaluationism and Barker’s version are susceptible to an objection from HIGHER-ORDER VAGUENESS. Even though these theories honor the intuition that there is no sharp boundary between the heavy things and the not-heavy things, they do entail that there is a sharp boundary between the heavy things and the borderline cases, and another sharp boundary between the borderline cases and the not-heavy things. However, many people believe that the meaning of *heavy* is also vague in this sense: there is no sharp boundary between the heavy things and the borderline cases, but only an imperceptible shading off from clear to less clear cases (e.g. Williamson, 1994, §5). We might try adding another layer of supervaluations, but then the argument could simply be re-run using the boundary between the clear cases of heavy things and the borderline borderline cases. Presumably the response would then be a further layer of supervaluations. Unless we can find a way to halt the regress, the supervaluational theory is in trouble.

The desire to avoid sharp boundaries at any level has led some theorists to advocate DEGREE-BASED theories. The classic treatment uses FUZZY LOGIC (Zadeh, 1978), according to which the classical truth-values are merely the extremes of a range  $[0, 1]$  of truth-values that sentences can take on. The slow fade from ‘heavy’ to ‘not heavy’ — the key feature which separates *heavy* from non-vague adjectives such as *geological* — is modeled by assuming a truth value-assigning function which does not have any sharp dis-

continuities. Theories of vagueness based on fuzzy logic have been advocated by a number of theorists including Goguen (1969); Lakoff (1973); Machina (1976); Zadeh (1978) and more recently Schiffer (2002); Smith (2010). There are many possible objections to this treatment, though: many theorists are unwilling to countenance degrees of truth on philosophical grounds, and there is good reason to think that fuzzy logic makes implausible predictions about the truth-values of compound sentences involving vague terms (Edgington, 1997). Furthermore, no such theory (to my knowledge) has given a satisfying story about where the function assigning truth-values to objects of a given measure should come from. It is also unclear how to account for the dependence of judgments about the applicability of vague terms on statistical facts about a reference class.

A rather different degree-based treatment is the STATISTICAL or PROBABILISTIC approach discussed by Borel (1907); Black (1937); Edgington (1997); Lawry (2008); Frazee & Beaver (2010); Lassiter (2011); Égré (2011); Lassiter & Goodman (2013); Égré & Barberousse (to appear). Glossing over some differences among these authors, the basic idea is that we treat vagueness without abandoning classical logic by treating the interpretation of *heavy* as statistical inference of the location of the unknown boundary  $\theta$ . Each possible value of  $\theta$  has some probability of being the true value, and the probability that an object counts as ‘heavy’ is simply the cumulative probability of  $\theta$  up to its weight:

$$\begin{aligned} P([x \text{ is heavy}]^{\mathcal{M},\theta} = 1) &= \int_0^\infty d\theta P([x \text{ is heavy}]^{\mathcal{M},\theta} = 1, \theta) \\ &= \int_0^\infty d\theta P([x \text{ is heavy}]^{\mathcal{M},\theta} = 1 | \theta) P(\theta) \\ &= \int_0^\infty d\theta P(\theta) \times \begin{cases} 1 & \text{if } \mathbf{weight}(x) \geq \theta \\ 0 & \text{otherwise} \end{cases} \\ &= \int_0^{\mathbf{weight}(x)} d\theta P(\theta) \end{aligned}$$

The derivation is determined by the mathematics of probability, except for the transition from the second to the third line, which derives from the classical semantics for *heavy* described in previous sections.

This approach allows us to reason about the relevant feature of the conversational context — the location of the boundary between *heavy* and *not heavy* — using the same machinery that we would use to infer the value of any other unknown variable for which probabilistic inference is appropriate. The probabilistic treatment is not truth-functional, in the sense that the probability of a conjunction or disjunction is not in general predictable from the probabilities of the conjuncts/disjuncts; it thus avoids several problems involving complex sentences which plague theories based on fuzzy logic (Edgington, 1997). This approach to vagueness can be given a precise compositional implementation

within a stochastic  $\lambda$ -calculus: see Lassiter & Goodman 2013 and chapter 24 of this handbook for details, and for a proposal to derive the context-sensitivity of relative adjectives and the specific form of the function  $P(\theta)$  from independently motivated devices: a probabilistic theory of uncertain reasoning and a coordination-based theory of interpretation.

## 4 Adjectives and scales

So far we have simply assumed that degrees are organized into ordered sets called ‘scales’ (or that the context-sensitive meanings of gradable adjectives display restrictions which mimic this treatment). Recently a good deal of attention has been devoted to detailed investigation of the different ways that scales can be organized, identifying several parameters of variation as well as different ways that the composition of scales can affect adjective interpretation.

### 4.1 Dimensionality

Dimensionality is relevant to scale structure in several distinct ways. First, we have cases in which adjectives are associated with distinct scales which support direct comparisons and share units, such as (32) (cf. Kennedy, 1997).

(32) The shelf is as tall as the table is wide.

The idea is that tallness and width can be compared in this way because their degree sets are the same. This is supported by the fact that both height and width can be measured in feet, inches, meters, etc. But *tall* and *wide* clearly differ in meaning: the former involves spatial extent in a vertical orientation, while the latter involves a horizontal direction. Following Kennedy (1997), then, we can suppose that scales are not simply composed of a set of degrees  $D$  and a partial order  $\geq$  on  $D$ , but also a dimension  $\delta$ : thus  $\langle D, \geq, \delta \rangle$ . The scales associated with *tall* and *wide* differ only in the dimension.

A second type of dimensionality effect involves the conceptual distinction between ONE-DIMENSIONAL adjectives such as *tall*, *wide*, and *heavy* and MULTIDIMENSIONAL adjectives such as *big*, *beautiful*, and *clever* (Bierwisch, 1989). For *heavy*, there is an unequivocal dimension along which individuals are measured: their weight. That is, suppose we know for certain exactly how much Sam and Bill weigh: then there can be no uncertainty about whether Sam is heavier than Bill or not, since weight is the only relevant dimension. But suppose Sam is taller and wider than Bill, but Bill is thicker and heavier. Which one is bigger? Even though these dimensions are all clearly relevant to the meaning of *big*, the fact that it is not immediately obvious which is bigger suggests that there is a certain indeterminacy in how height, width, thickness, etc. are taken into account in determining the ordering of objects in terms of their ‘bigness’.

Note, however, that it is possible to use a multidimensional adjective such as *big* with an explicit specification of the relevant dimension: both of the sentences in (33) would be true in the context described.

- (33) a. Sam is bigger than Bill with respect to height.  
 b. Bill is bigger than Sam with respect to weight.

In this case, it appears that the explicit specification of a dimension serves to narrow temporarily the set of dimensions relevant of the interpretation of the adjective.

What about cases in which no dimension is specified – what does it take for Sam to be bigger than Bill *simpliciter*? For simplicity, let’s assume that there is a finite number of discrete dimensions that are relevant to the meaning of any adjective. Suppose now that we have an exhaustive list  $\Delta = \{\delta_1, \delta_2, \dots, \delta_n\}$  of the  $n$  dimensions that are relevant to the meaning of *big*. There are at least three possibilities. First, we might suppose that the context of utterance supplies a dimension. This seems wrong, though: *bigger* is not simply ambiguous between meaning ‘heavier’, ‘wider’, ‘taller’, etc., but rather seems to take into account information about all of these dimensions simultaneously. A second possibility is that multidimensional adjectives universally quantify over relevant dimensions. On this theory, Sam is bigger than Bill if and only if, for all  $i \in \{1, 2, \dots, n\}$ , Sam is bigger than Bill if *big* is interpreted with respect to the scale  $\langle D, \geq, \delta_i \rangle$ . This type of analysis would make the BIGNESS scale non-connected, since neither ‘Sam is bigger than Bill’ nor ‘Bill is bigger than Sam’ is true if Sam is bigger on some dimensions and Bill is bigger on others. It could also be that multidimensional adjectives vary in what type of quantification over dimensions they invoke. Yoon (1996); Sassoon (2013) argue that this is the case for *healthy* and *sick*: to be healthy you have to be healthy in *every* way, but you are sick if you are sick in even *one* way, even if you are healthy in all others.

A somewhat different idea is that the scales associated with multidimensional adjectives are constructed using a context-sensitive function which collapses objects’ measurements along the various relevant dimensions into a single scale, taking into account information about all relevant dimensions but possibly weighting them differently. While this type of construct is less familiar to linguists, psychologists interested in how people map high-dimensional spaces to low-dimensional ones have investigated a number of such techniques (see e.g. Markman, 1998, §2 for an overview and pointers to the extensive literature on relevant topics from psychology).

## 4.2 Antonymy

Natural language adjectives frequently come in pairs of ANTONYMS: some uncontroversial examples are *tall/short*, *dangerous/safe*, *full/empty*, *heavy/light*, and *early/late*. As I’ll use it,  $A_1$  and  $A_2$  are antonyms if and only if the following is trivially true.

(34) For all  $x$  and  $y$ :  $x$  is more  $A_1$  than  $y$  if and only if  $y$  is more  $A_2$  than  $x$ .

Given this characterization, it is reasonable to assume that antonymous pairs of adjectives are adjectives which share a set of degrees and a dimension, but differ in that the ordering is reversed. For example,  $x$  is heavier than  $y$  if and only if  $x$ ’s weight is *greater* than  $y$ ’s, and  $x$  is lighter than  $y$  if and only if  $x$ ’s

weight is *less* than *y*'s. More generally, we can stipulate that if  $A_1$  and  $A_2$  are antonyms and  $A_1$  is lexically associated with the scale  $\langle D, \geq, \delta \rangle$ , then  $A_2$  is lexically associated with the scale  $\langle D, \leq, \delta \rangle$ , where  $\leq = \geq^1 = \{(d, d') \mid (d', d) \in \geq\}$ . (34) then follows.

There is an interesting subtlety, though: as a rule, one of the members of a pair is 'marked' or 'evaluative' in the sense that the comparative strongly implies that the corresponding positive form holds of one of the members. I'll call the member of the pair without this property its 'positive' member, and the marked member its 'negative' member. Note first that the sentence in (35a) is unremarkable with the positive adjective *heavy*, indicating that the inference in (35b) is not a good one.

- (35) a. Box *A* is heavier than box *B*, but both are quite light.  
 b. Box *A* is heavier than box *B*.  $\not\rightarrow$  Box *A* is heavy.

Other positive adjectives are similar. However, it has been argued that (36a) is less acceptable, and that it is natural to draw inferences with the form of (36b) for negative adjectives such as *light*.

- (36) a. ? Box *A* is lighter than box *B*, but both are quite heavy.  
 b. Box *A* is lighter than box *B*.  $\sim$  Box *A* is light.

(36b) is at best a pragmatic inference, given that (36a) is not an outright contradiction. Indeed, it is not difficult to find naturally-occurring examples similar to (36a): (37) gives two found on the web.

- (37) a. Tried putting her on her side in the mud and pulling and pushing her to get her out of the rut. No luck, she may be lighter than Maria but still damn heavy. [DL note: 'Maria' and 'her' are motorcycles.]  
 b. I have a Britax Roundabout, which is slightly smaller than the Marathon, but it's really big!

Why the use of a negative adjective should license such an inference in some circumstances — however weakly — is not entirely clear: see Bierwisch 1989; Rett 2008a; Sassoon 2010b for discussion and further references.

### 4.3 Adjective type, boundedness, and degree modification

As Hay *et al.* (1999); Rotstein & Winter (2004); Kennedy & McNally (2005); Kennedy (2007) discuss, scales could logically come in any of four types with respect to their BOUNDEDNESS. They can either have or lack a unique greatest element (a maximum), and they can either have or lack a unique least element (a minimum). Scales with neither a maximum nor a minimum element are FULLY OPEN; those with both are FULLY CLOSED. Scales with a minimum but no maximum are LOWER CLOSED, and those with a maximum but no minimum are UPPER CLOSED. The latter two types are of course formally identical except for the choice of the default/unmarked polarity. (See Jackendoff 1991; Paradis 2001 for additional relevant considerations.)

These are merely logical possibilities, but Kennedy & McNally (2005) argue that all four scale types are attested in English. The following four adjective pairs can be used to illustrate the proposal:

- (38) a. Fully closed: *empty/full*  
 b. Fully open: *ugly/beautiful*  
 c. Upper closed: *impure/pure*  
 d. Lower closed: *straight/bent*

*Full* and *empty* are antonyms and are both intuitively associated with maxima — that is, there is a principled limit to how full or empty something could get, i.e. a point after which you could not make the object any more full/empty. Corroborating this intuition are a number of linguistic tests, of which we discuss two here. First, these adjectives in the positive form seem to associate with the maximum point. That is, if someone tells you that a theater is full (empty) you expect that there are no or almost no empty (full) seats in it. It does not, for example, mean only that the theater is more full (more empty) than normal. These are thus both examples of MAXIMUM-STANDARD adjectives. Second, both of these adjectives can be modified by *completely*, *perfectly*, and *maximally*, and the result reinforces adherence to a maximum point: *The theater is completely full* means that you will have to go elsewhere to watch the movie. If ‘completely/perfectly/maximally *A*’ means ‘having the maximum possible degree of the scalar property *A*’, then the acceptability of this collocation in this meaning indicates that the scale in question has a maximum possible degree. Finally, Kennedy (2007) points out that the meanings of adjectives like *full* and *empty* are much less uncertain than those of prototypical vague adjectives, and that the sorites paradox (32) is less compelling with them. Kennedy argues that this is explained if the meanings of these adjectives are associated with the relevant scalar endpoints, rather than being fixed by contextual information. If this is all correct, then the scale associated with the adjective pair *empty/full* must have a maximum and a minimum point, and so is fully closed.

*Ugly* and *beautiful*, on the other hand, are prototypical vague adjectives of the type discussed above. It also seems clear that neither is associated with a maximum — that is, there is no principled limit to how beautiful or ugly something could be. Corroborating this intuition, these adjectives are rather odd with *perfectly* and *maximally*. (They are acceptable with *completely*, but with a different meaning that does not seem to be degree-modifying.) The lack of endpoints forces these adjectives to be RELATIVE-STANDARD, a fact which Kennedy (2007) argues to be connected to their vagueness.

The pair *impure/pure* in (38c) is different in that one member appears to make reference to a maximal degree, and the other seems to invoke deviation from that degree in the opposite direction. That is, if a sample is pure it could not be more pure; but a sample could be impure even when there are ways to make it still more impure. Corroborating these intuitions, *completely/perfectly/maximally pure* is acceptable while the same modifiers are

rather off with *impure*. *Pure* is thus a maximum-standard adjective, while *impure* is a MINIMUM-STANDARD adjective, indicating deviation from complete purity, however small. According to Kennedy (2007), these adjectives are also less vague than relative adjectives. Similar considerations hold of the *straight/bent* pair in (38d), with *straight* corresponding to *pure* and *bent* to *impure*.

Two questions suggest themselves. First, are there any relative adjectives which fall on scales with minimum and/or maximum points — that is, can the meanings of positive-form adjectives fail to be ‘attracted’ to the endpoints when there are endpoints present? Second, if minimum- and maximum-standard adjectives are not vague, how can this fact be explained within the context of a general theory of vagueness of the type described in §3.4? Kennedy (2007) argues that the answer to the first question is ‘no’, and suggests a pragmatic/processing principle of ‘Interpretive Economy’ designed to explain this gap while also answering the second question. Potts (2008) criticizes this account and proposes a derivation of Interpretive Economy as a historical tendency from a game-theoretic perspective on communication.

However, Lassiter (2010a) and McNally (2011) have pointed out several apparent examples of relative adjectives which fall onto non-open scales, including *expensive/inexpensive*, *likely/unlikely*, *probable/improbable*, and relative uses of *full*. If correct, these data would falsify the categorical empirical claim motivating both Kennedy’s and Potts’ accounts. Lassiter & Goodman (2013) propose a probabilistic coordination-based account which suggests an explanation of the correlation between adjective meanings and scale structure, but also allows for deviations under specific circumstances.

## 5 Comparatives and degree operator scope

### 5.1 A theory of comparatives

This section sketches briefly a treatment of comparatives in a degree-based theory, referring readers to literature cited for many more details and alternatives. See also Schwarzschild (2008) for an introductory essay on comparatives, and von Stechow (1984); Klein (1991); Beck (2012) for more detailed surveys.

COMPARATIVE sentences relate the measures of two objects along a scale or, less often, along two different scales. The simplest case involves sentences such as (39).

- (39) Sam is taller than Bill is.  
 ‘Sam’s height is greater than Bill’s height.’

We will assume that there is ellipsis in the comparative clause in (39), so that it is interpreted as if it were *Sam is taller than Bill is tall*.

EQUATIVE sentences are closely related in meaning to comparatives: it is usually assumed that they are related by a simple change from ‘>’ to ‘≥’ in the definitions. Equatives are much less studied than comparatives, though, and there may well be interesting differences between them.

- (40) Sam is as tall as Bill is.  
 ‘Sam’s height is (at least) as great as Bill’s height.’

An influential treatment of comparatives associated with von Stechow (1984) starts with the treatment of gradable adjectives as functions from a degree to a property, as in (42).

$$(41) \llbracket \text{tall} \rrbracket^{\mathcal{M}} = \lambda d_d \lambda x_e [\mathbf{height}(x) \geq d]$$

This is an ‘at least’ meaning: *x is 5 feet tall* will come out as true of any *x* whose height is 5 feet or greater. (The fact that *Sam is 5’ tall* seems odd if he is in fact 6’ tall can be explained as an effect of a quantity implicature typically associated with an assertion of this sentence, to the effect that Sam’s height does not exceed 5’.) Now, an initial attempt to state truth-conditions for (39) runs like this: there are degrees of height  $d_1$  and  $d_2$  such that  $\mathbf{height}(\mathbf{Sam}) \geq d_1$  and  $\mathbf{height}(\mathbf{Bill}) \geq d_2$  and  $d_1 > d_2$ . The problem is that these truth-conditions are true even if Sam is *shorter* than Bill. To see this, suppose that Sam is 5’ tall and Bill is 6’ tall, and set  $d_1 = 4'$  and  $d_2 = 3'$ . The three conditions are satisfied, since Sam’s height is greater than or equal to  $d_1$  and Bill’s height is greater than or equal to  $d_2$ . But this is clearly not a situation which verifies *Sam is taller than Bill is*, and so these are the wrong truth-conditions.

The solution is to consider the *maximal* degree of height that Sam and Bill have. That is, we define an operator **max** which returns the greatest member of a set of degrees.

$$(42) \mathbf{max} = \lambda D_{\langle d, t \rangle} \iota d_d [D(d) \wedge \forall d' [D(d') \rightarrow d \geq d']]$$

We then define the comparative morpheme *more/-er* as a function which takes two degree sets as input and compares their respective maxima.

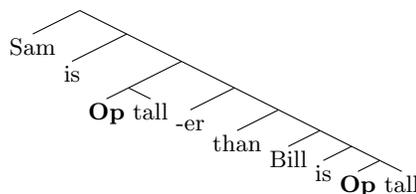
$$(43) [|\textit{more/-er}|]^{\mathcal{M}} = \lambda D_{\langle d,t \rangle} \lambda D'_{\langle d,t \rangle} [\mathbf{max}(D') > \mathbf{max}(D)]$$

To make this denotation produce the right result, we have to ensure that the first argument is (the characteristic function of) the set of degrees to which Bill is tall, and the second argument is (the characteristic function of) the set of degrees to which Sam is tall. If so, the sentence will return the value True if and only if Sam's greatest degree of height is greater than Bill's, i.e., if Sam is taller:

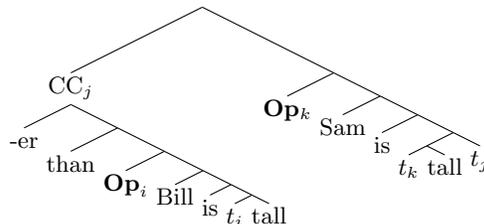
$$(44) [|\textit{Sam is taller than Bill is}|]^{\mathcal{M}} = \\
\begin{aligned}
& [|\textit{-er}|]^{\mathcal{M}} (\lambda d_d [\mathbf{height}(\mathbf{Bill}) \geq d]) (\lambda d_d [\mathbf{height}(\mathbf{Sam}) \geq d]) = \\
& \mathbf{max}(\lambda d_d [\mathbf{height}(\mathbf{Sam}) \geq d]) > \mathbf{max}(\lambda d_d [\mathbf{height}(\mathbf{Bill}) \geq d]) = \\
& = \mathbf{height}(\mathbf{Sam}) > \mathbf{height}(\mathbf{Bill})
\end{aligned}$$

It is a non-trivial matter to engineer a Logical Form which has this property in a syntactically responsible way, though. It requires us to assume that the comparative clause is a complex scope-taking expression which undergoes Quantifier Raising and the presence of silent operators whose movement triggers  $\lambda$ -abstraction of a degree variable (or something equivalent in semantic effect). The following trees depict one surface structure (top) and logical form (bottom) which would generate the right truth-conditions when combined with our other assumptions.

**Surface structure:**



**Logical form:**



**Op** movement triggers  $\lambda$ -abstraction of a degree variable in each clause, and the whole comparative clause must undergo Quantifier Raising. Assuming that *than* is semantically vacuous, *-er* will now combine first with the clause [**Op<sub>i</sub>** *Sam is t<sub>i</sub> tall*], and then with the clause [**Op<sub>k</sub>** *Sam is t<sub>k</sub> tall*]. The reader may check that this LF derives the truth-conditions spelled out in

(44). (See Heim (2001); Lassiter (2013) for further discussion of this particular nexus of assumptions.)

## 5.2 Scope interactions between degree operators, modals, and quantifiers

There are, to be sure, many possible alternative ways to derive the truth-conditions in (44), some of which are considerably less complex than the derivation just presented. For example, Kennedy (1997) presents a syntax and semantics in which the comparative clause denotes a degree (not a set of degrees) and is interpreted without LF-movement. What is worse, it is not hard to see that the result of combining of the comparative’s **max** operator (42) with a meaning like  $\lambda d_d[\mathbf{height}(\mathbf{Sam}) \geq d]$  will always be equivalent to the much simpler **height(Sam)**. Why the extra complications?

The motivation for this roundabout way of calculating truth-conditions for *Sam is taller than Bill is* is that it makes room for scope interactions with quantificational and modal elements. Consider (45), based on an example from Heim (2001).

(45) Iowa City is closer to Lake Michigan than it is to an ocean.

This sentence can be read in two ways, but the intended meaning is that the *closest* ocean to Iowa City is not as close as Lake Michigan — that is, that there is not any ocean which is closer to Iowa City than Lake Michigan. Once we have a quantificational element involved (the existential *an ocean*), our roundabout way of calculating truth-conditions comes in handy: our treatment predicts (46) as one possible reading of (45). (Note that *close* takes one degree and two individual arguments, since it expresses the degree of closeness between two locations.)

(46)  $[[[(45)]]^M = \mathbf{close}(\mathbf{IC})(\mathbf{LM}) > \mathbf{max}(\lambda d_d \exists x_e [\mathbf{ocean}(x) \wedge \mathbf{close}(\mathbf{IC})(x) \geq d])]$

“The closeness between LM and IC is greater than the greatest degree of closeness s.t. there is an ocean that close to IC; i.e., the distance is less than the distance to the closest ocean”

Similarly, modals such as *allowed*, *have to*, and *required* interact scopally with comparatives. Imagine (47) spoken by an amusement park employee to a disappointed child.

(47) You’re 4’ tall; you have to be exactly 1’ taller than that in order to ride on this ride.

This sentence can be read in two ways. The implausible reading would entail that only people who are *exactly* 5’ tall can ride on the ride. A more plausible interpretation is that the child is exactly 1’ too short to ride, i.e. the requirement is that riders be at *at least* 5’ tall.

To show how this semantics derives both, we assume that *have to* denotes a universal quantifier over some set **Acc** of accessible worlds, and we intensionalize the interpretation of the adjective (i.e., we add a world argument so that the adjective is of type  $\langle s, \langle d, et \rangle \rangle$ ). The comparative clause *exactly 1' -er than that* enforces equality with 5' (by a straightforward compositional process that we won't pause to spell out here). We also use **a** to rigidly designate the addressee. We can now generate both readings by varying the scope of the modal and the comparative clause:

(48) **Reading 1:** *have to* > *CC*  
 $\forall w \in \mathbf{Acc} : \mathbf{max}(\lambda d[\mathbf{height}(w)(\mathbf{a}) \geq d]) = 5'$   
 'The addressee is exactly 5' tall in all accessible riding-worlds'

(49) **Reading 2:** *CC* > *have to*  
 $\mathbf{max}(\lambda d[\forall w \in \mathbf{Acc} : \mathbf{height}(w)(\mathbf{a}) \geq d]) = 5'$   
 'The greatest degree *d* such that the addressee is *at least d*-tall in all accessible riding-worlds is 5' — that is, in all accessible riding-worlds the addressee is at least 5' tall'

There is a still-unresolved problem here, though. Heim (2001) points out that the critical assumptions allowing our semantics to generate this ambiguity also predict that a similar ambiguity should appear with a universal DP such as *everyone who rode*. But the corresponding sentences are clearly not ambiguous:

(50) This kid is 4' tall. Everyone who rode was exactly 1' taller than that.

(51) a. **Attested:** 'Every rider was exactly 5' tall'  
 b. **Unattested:** 'Every rider was at least 5' tall'

(51b) is not a possible reading of (50): that is, (50) is false if any rider was taller than 5'. This is puzzling, since it is standardly assumed that strong modals like *have to* and quantifier phrases such as *everyone* differ only in that the former quantify over worlds and the latter over individuals, and we expect to find the same scope ambiguities with both — including a reading that matches the 'at least' meaning of the sentence with *have to*. What is even more puzzling is that the same ambiguities do appear sporadically with existentially quantified DPs, as in our example (45) above.

The pattern of attested and missing readings is not specific to comparative constructions: Szabolcsi (2006); Lassiter (2010b, 2013) point out that it has a precise parallel in the pattern of restrictions on the interactions between modals/quantifiers and amount *wh*-questions treated in the literature under the name of WEAK ISLANDS. While there is no fully worked-out and agreed-upon theory of weak islands involving amount *wh*-expressions, there is some work which suggests an explanation of the divergence between existential and universal DPs: Szabolcsi & Zwarts (1993) argue that certain semantic operations

required to compute the effect of universal quantification are undefined in the domain of degrees.

If this account is correct, though, it should also affect modals which are modeled as universal quantifiers over worlds, contrary to fact: comparatives with strong modals are sometimes ambiguous where comparatives with universally quantified DPs are not (compare (47)-(49) to (50)-(51)). Lassiter (2013) uses this divergence to argue that strong modals are not in fact universal quantifiers over worlds, but rather degree expressions which take propositional arguments. On this account, the scope interactions in (47)-(49) are not between a degree operator and a quantifier over worlds, but between two different kinds of degree operators. The semantic restrictions on the operation of universal quantification suggested by Szabolcsi & Zwarts (1993) would thus not apply, since neither reading of (47) makes reference to universal quantification over worlds.

Whether or not this account of the detailed patterns of scope interactions is ultimately successful, there is a wide variety of independent arguments which provide general motivation for a degree-based treatment of modality over a quantificational one: see Lassiter 2014 for an extended treatment.

## 6 Conclusion

Understanding adjectival meaning requires simultaneous attention to the morphology, syntax, semantics, and pragmatics of modification and gradation. This chapter has presented a selective overview of phenomena and puzzles involving adjectival modification and gradation: the semantics of adjective-noun modification constructions; morphological, syntactic, and semantic issues involving gradation, and degrees; the compositional semantics and pragmatics of vague adjectives; scale structure and other typological distinctions among gradable adjectives; and some puzzles in the interaction between comparatives and quantifiers.

Much more remains to be said, for example, about degree modification (Bolinger, 1972; Klein, 1980; Rotstein & Winter, 2004; Kennedy & McNally, 2005; Kennedy, 2007; Rett, 2008b; McNabb, 2012; Morzycki, 2013) and superlatives (Heim, 1985; Szabolcsi, 1986; Gawron, 1995; Farkas & Kiss, 2000; Sharvit & Stateva, 2002; Teodorescu, 2009; Bobaljik, 2012; Szabolcsi, 2013). Hopefully this chapter will have provided sufficient background for the reader to dig deeper into this rich and fascinating literature.

## References

- Barker, Chris (2002), The dynamics of vagueness, *Linguistics and Philosophy* 25(1):1–36.
- Barker, Chris & Chung-Chieh Shan (2014), *Continuations and natural language*, Oxford University Press.
- Bartsch, Renate & Theo Vennemann (1973), *Semantic Structures: A Study In the Relation Between Semantics and Syntax*, Athenäum.
- Beck, Sigrid (2012), Comparison constructions, in Claudia Maienborn, Klaus von Heusinger, & Paul Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning*, Mouton de Gruyter, volume 2, (1341–1390).
- Beck, Sigrid, Sveta Krasikova, Daniel Fleischer, Remus Gergel, Stefan Hofstetter, Christiane Savelsberg, John Vanderelst, & Elisabeth Villalta (2009), Cross-linguistic variation in comparison constructions, *Linguistic Variation Yearbook* 9(1):1–66.
- van Benthem, Johan (1983), *The Logic of Time*, Reidel.
- Bierwisch, Manfred (1989), The semantics of gradation, in M. Bierwisch & E. Lang (eds.), *Dimensional Adjectives: Grammatical Structure and Conceptual Interpretation*, Springer-Verlag, (71–261).
- Black, Max (1937), Vagueness. An exercise in logical analysis, *Philosophy of Science* 4(4):427–455.
- Bobaljik, Jonathan D. (2012), *Universals in Comparative Morphology: Suppletion, superlatives, and the structure of words*, MIT Press.
- Bochnak, M. Ryan (2013), *Cross-linguistic variation in the semantics of comparatives*, Ph.D. thesis.
- Bolinger, Dwight (1972), *Degree words*, Mouton.
- Borel, Émile (1907), Sur un paradoxe économique: Le sophisme du tas de blé et les vérités statistiques, *Revue du Mois* .
- Burnett, Heather (2012), The puzzle(s) of absolute adjectives: On vagueness, comparison, and the origin of scale structure, *UCLA Working Papers in Linguistics* 16:1–50.
- Chomsky, Noam (1986), *Knowledge of language: Its nature, origin, and use*, Praeger.
- Cresswell, Max J. (1976), The semantics of degree, in B. Partee (ed.), *Montague grammar*, Academic Press, (261–292).
- Cresswell, Max J. (1990), *Entities and indices*, Kluwer.
- Davidson, Donald (1967), The logical form of action sentences, *The logic of decision and action* .
- van Deemter, Kees (2010), *Not exactly: In praise of vagueness*, Oxford University Press.
- Doetjes, Jenny, Camelia Constantinescu, et al. (2009), A neo-kleinian approach to comparatives, in *Semantics and Linguistic Theory (SALT) 19*, (124–141).
- Edgington, Dorothy (1997), Vagueness by degrees, in Rosanna Keefe & Peter Smith (eds.), *Vagueness: A Reader*, MIT Press, (294–316).
- Égré, Paul (2011), Perceptual ambiguity and the sorites, in Rick Nouwen, Robert van Rooij, Uli Sauerland, & Hans-Christian Schmitz (eds.), *Vagueness in Communication*, Springer, (64–90).
- Égré, Paul & Anouk Barberousse (to appear), Borel on the heap, *Erkenntnis* .
- Fara, Delia Graff (2000), Shifting sands: An interest-relative theory of vagueness, *Philosophical Topics* 20:45–81.

- Farkas, Donka F & Katalin É Kiss (2000), On the comparative and absolute readings of superlatives, *Natural Language & Linguistic Theory* 18(3):417–455.
- Fine, Kit (1975), Vagueness, truth and logic, *Synthese* 30(3):265–300.
- Fraeeze, Joey & David Beaver (2010), Vagueness is rational under uncertainty, in Maria Aloni, Harald Bastiaanse, Tikitou de Jager, & Katrin Schulz (eds.), *Logic, Language and Meaning: 17th Amsterdam Colloquium, Amsterdam, The Netherlands, December 16-18, 2009, Revised Selected Papers*, Springer, (153–162).
- Gallin, Daniel (1975), *Intensional and higher-order modal logic: With applications to Montague semantics*, Elsevier.
- Gawron, Jean Mark (1995), Comparatives, superlatives, and resolution, *Linguistics and Philosophy* 18(4):333–380.
- Goguen, Joseph A (1969), The logic of inexact concepts, *Synthese* 19(3):325–373.
- Grano, Thomas (2012), Mandarin *hen* and universal markedness in gradable adjectives, *Natural Language & Linguistic Theory* :1–53.
- Halle, Morris & Alec Marantz (1993), Distributed morphology and the pieces of inflection, in Kenneth L. Hale & Samuel J. Keyser (eds.), *The view from Building 20*, MIT Press, (111–176).
- Hay, Jennifer, Chris Kennedy, & Beth Levin (1999), Scalar structure underlies telicity in degree achievements, in *Semantics and Linguistic Theory 9*, (127–144).
- Heim, Irene (1985), Notes on comparatives and related matters, Ms., University of Texas, Austin .
- Heim, Irene (2001), Degree operators and scope, in Fery & Sternefeld (eds.), *Audiatu Vox Sapientiae: A Festschrift for Arnim von Stechow*, Akademie Verlag, (214–239).
- Huddleston, Rodney & Geoffrey K. Pullum (2002), *The Cambridge Grammar of English*, Cambridge University Press.
- Jackendoff, Ray (1991), Parts and boundaries, *Cognition* 41(1):9–45.
- Jacobson, Pauline (1999), Towards a variable-free semantics, *Linguistics and Philosophy* 22(2):117–185.
- Kamp, Hans (1975), Two theories about adjectives, in E. Keenan (ed.), *Formal Semantics of Natural Language*, Cambridge University Press, (123–155).
- Kamp, Hans & Barbara Partee (1995), Prototype theory and compositionality, *Cognition* 57(2):129–191.
- Keefe, Rosanne (2000), *Theories of Vagueness*, Cambridge University Press.
- Keefe, Rosanne & Peter J. Smith (eds.) (1997), *Vagueness: A Reader*, MIT Press.
- Kennedy, Chris (1997), *Projecting the adjective: The syntax and semantics of gradability and comparison*, Ph.D. thesis, U.C., Santa Cruz.
- Kennedy, Chris (2001), Polar opposition and the ontology of ‘degrees’, *Linguistics and Philosophy* 24(1):33–70.
- Kennedy, Chris (2007), Vagueness and grammar: The semantics of relative and absolute gradable adjectives, *Linguistics and Philosophy* 30(1):1–45.
- Kennedy, Chris & Louise McNally (2005), Scale structure, degree modification, and the semantics of gradable predicates, *Language* 81(2):345–381.
- Klein, Ewan (1980), A semantics for positive and comparative adjectives, *Linguistics and Philosophy* 4(1):1–45.
- Klein, Ewan (1991), Comparatives, in A. von Stechow & D. Wunderlich (eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, Walter de Gruyter.

- Krantz, David H., R. Duncan Luce, Patrick Suppes, & Amos Tversky (1971), *Foundations of Measurement*, Academic Press.
- Lakoff, George (1973), Hedges: A study in meaning criteria and the logic of fuzzy concepts, *Journal of philosophical logic* 2(4):458–508.
- Larson, Richard (1998), Events and modification in nominals, in *Proceedings of SALT*, volume 8, (145–168).
- Lassiter, Daniel (2010a), Gradable epistemic modals, probability, and scale structure, in Nan Li & David Lutz (eds.), *Semantics & Linguistic Theory (SALT) 20*, CLC Publications, (197–215).
- Lassiter, Daniel (2010b), The Algebraic Structure of Amounts: Evidence from Comparatives, *Interfaces: Explorations in Logic, Language and Computation* :38–56.
- Lassiter, Daniel (2011), Vagueness as probabilistic linguistic knowledge, in R. Nouwen, R. van Rooij, U. Sauerland, & H.-C. Schmitz (eds.), *Vagueness in Communication*, Springer, (127–150).
- Lassiter, Daniel (2013), Quantificational and modal interveners in degree constructions, in A. Chereches (ed.), *Semantics & Linguistic Theory (SALT) 22*, CLC Publications.
- Lassiter, Daniel (2014), *Measurement and Modality: The Scalar Basis of Modal Semantics*, Oxford University Press (to appear).
- Lassiter, Daniel & Noah D. Goodman (2013), Context, scale structure, and statistics in the interpretation of positive-form adjectives, to appear in *Semantics & Linguistic Theory (SALT) 23*.
- Lawry, Jonathan (2008), Appropriateness measures: an uncertainty model for vague concepts, *Synthese* 161(2):255–269.
- Lewis, David (1970), General semantics, *Synthese* 22(1):18–67.
- Liu, Chen-Sheng Luther (2010), The positive morpheme in chinese and the adjectival structure, *Lingua* 120(4):1010–1056.
- Machina, Kenton F (1976), Truth, belief, and vagueness, *Journal of Philosophical logic* 5(1):47–78.
- Markman, Arthur B (1998), *Knowledge Representation*, Lawrence Erlbaum.
- McNabb, Yaron (2012), *The Syntax and Semantics of Degree Modification*, Ph.D. thesis, University of Chicago.
- McNally, Louise (2011), The relative role of property type and scale structure in explaining the behavior of gradable adjectives, in Rick Nouwen, Robert van Rooij, Uli Sauerland, & Hans-Christian Schmitz (eds.), *Vagueness in Communication*, Springer, (151–168).
- Montague, Richard (1970), English as a formal language, *Linguaggi nella società e nella Tecnica* :189–224.
- Montague, Richard (1973), The proper treatment of quantification in ordinary English, in J. Hintikka, J. Moravcsik, & P. Suppes (eds.), *Approaches to Natural Language*, Reidel, volume 49, (221–242).
- Morzycki, Marcin (2013), *Modification*, Cambridge University Press, to appear.
- Murphy, Gregory (2002), *The Big Book of Concepts*, MIT Press.
- Paradis, Carita (2001), Adjectives and boundedness, *Cognitive Linguistics* 12(1):47–66.
- Parsons, Terence (1990), *Events in the Semantics of English*, MIT Press.
- Partee, Barbara (1995), Lexical semantics and compositionality, in Lila Gleitman & Mark Liberman (eds.), *An invitation to cognitive science*, volume 1, (311–360).

- Partee, Barbara H (2007), Compositionality and coercion in semantics: The dynamics of adjective meaning, in Gerlof Bouma, Irene Krämer, & Joost Zwarts (eds.), *Cognitive foundations of interpretation*, University of Chicago Press, (145–161).
- Partee, Barbara H., Alice ter Meulen, & Robert E. Wall (1990), *Mathematical methods in linguistics*, Springer.
- Potts, Christopher (2008), Interpretive Economy, Schelling points, and evolutionary stability, ms., University of Massachusetts at Amherst.
- Rett, Jessica (2008a), Antonymy and Evaluativity, in T. Friedman & M. Gibson (eds.), *Semantics and Linguistic Theory (SALT) 17*, CLC Publications, Ithaca, NY, (210–227).
- Rett, Jessica (2008b), *Degree modification in natural language*, Ph.D. thesis, Rutgers.
- van Rooij, Robert (2011a), Measurement, and interadjective comparisons, *Journal of Semantics* 28(3):335–358.
- van Rooij, Robert (2011b), Vagueness and linguistics, in G. Ronzitti (ed.), *Vagueness: A guide*, Springer, (123–170).
- Rotstein, Carmen & Yoad Winter (2004), Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers, *Natural Language Semantics* 12(3):259–288.
- Rullmann, Hotze (1995), *Maximality in the Semantics of wh-constructions*, Ph.D. thesis, University of Massachusetts, Amherst.
- Sassoon, Galit (2010a), Measurement theory in linguistics, *Synthese* 174(1):151–180, ISSN 0039-7857.
- Sassoon, Galit (2010b), The degree functions of negative adjectives, *Natural Language Semantics* 18(2):141–181.
- Sassoon, Galit W. (2013), A typology of multidimensional adjectives, *Journal of Semantics* 30(3):335–380.
- Schiffer, Stephen (2002), Vagueness and Partial Belief, *Noûs* 34(1):220–257.
- Schwarzschild, Roger (2008), The semantics of comparatives and other degree constructions, *Language and Linguistics Compass* 2(2):308–331.
- Schwarzschild, Roger & Karina Wilkinson (2002), Quantifiers in comparatives: A semantics of degree based on intervals, *Natural Language Semantics* 10(1):1–41.
- Shapiro, Stewart (2006), *Vagueness in context*, Clarendon Press.
- Sharvit, Yael & Penka Stateva (2002), Superlative expressions, context, and focus, *Linguistics and philosophy* 25(4):453–504.
- Siegel, Muffy Ann (1976), *Capturing the Adjective*, Ph.D. thesis, University of Massachusetts at Amherst.
- Smith, Nicholas J. J. (2010), Degree of belief is expected truth value, in R. Dietz & S. Moruzzi (eds.), *Cuts and Clouds: Vagueness, Its Nature and Its Logic*, Oxford University Press.
- Solt, Stephanie (2009), *The semantics of adjectives of quantity*, Ph.D. thesis, City University of New York.
- Solt, Stephanie (2011), Notes on the comparison class, in Rick Nouwen, Robert van Rooij, Uli Sauerland, & Hans-Christian Schmitz (eds.), *Vagueness in Communication*, Springer, (127–150).
- von Stechow, Arnim (1984), Comparing semantic theories of comparison, *Journal of Semantics* 3(1):1–77.
- Steedman, M. (1987), Combinatory grammars and parasitic gaps, *Natural Language & Linguistic Theory* 5(3):403–439.
- Steedman, Mark (2001), *The Syntactic Process*, MIT press.

- Szabolcsi, Anna (1986), Comparative superlatives, *MIT Working papers in Linguistics* 8:245–265.
- Szabolcsi, Anna (1987), Bound variables in syntax: Are there any?, in J. Groenendijk et al. (ed.), *Proceedings of the 6th Amsterdam Colloquium*, (331–353).
- Szabolcsi, Anna (2006), Strong vs. weak islands, *Blackwell Companion to Syntax* 4:479–531.
- Szabolcsi, Anna (2013), Compositionality without word boundaries: (the) more and (the) most, in *Semantics & Linguistic Theory (SALT) 22*, (1–25).
- Szabolcsi, Anna & Frans Zwarts (1993), Weak islands and an algebraic semantics for scope taking, *Natural Language Semantics* 1(3):235–284.
- Teodorescu, Viorica Alexandra (2009), *Modification in the noun phrase: the syntax, semantics, and pragmatics of adjectives and superlatives*, Ph.D. thesis, University of Texas at Austin.
- Wellwood, Alexis, Valentine Hacquard, & Roumyana Pancheva (2012), Measuring and comparing individuals and events, *Journal of semantics* 29(2):207–228.
- Williamson, Timothy (1994), *Vagueness*, Routledge.
- Wright, Crispin (1976), Language-mastery and the sorites paradox, in Gareth Evans & John McDowell (eds.), *Truth and Meaning*, Oxford: Clarendon University Press, (223–47).
- Yoon, Youngeun (1996), Total and partial predicates and the weak and strong interpretations, *Natural language semantics* 4(3):217–236.
- Zadeh, Lofti A. (1978), Fuzzy sets as a basis for a theory of possibility, *Fuzzy Sets and Systems* 1(1):3–28, ISSN 0165-0114.