The Sorites Paradox

Sergi Oms and Elia Zardini

Contributors

1

The Sorites Paradox in Linguistics

Christopher Kennedy

1.1 Introduction

Statements of the Sorites Paradox often take the form in (1), where

(1) If x is α and there is a y such that x is just a bit more α than y, then y is α.

More often than not, α is a gradable adjective such as *long, tall, short, young, happy, healthy, tasty* or *bald*, which can be substituted directly for α in (1) without changing anything else about the sentence. But with slight modifications to the surrounding syntax, α can take the form of other grammatical categories as well: it can be a noun or noun phrase (*if x is a heap of sand and there is a y such that x is just a bit more of a heap of sand than y, then y is a heap of sand*), a verb or verb phrase (*if x likes broccoli and there is a y such that x likes broccoli just a bit more than y, then y likes broccoli*), a prepositional phrase (*if x is near the ocean and there is a y such that x is just a bit nearer the ocean than y, then y is near the ocean*), and so forth.

These kinds of examples illustrate the feature of vague predicates that has been of primary interest to linguists: independent of grammatical category, vague predicates generally have both an unmarked "positive" form — the bare version of α that appears both in the initial premise of the argument and in the apparently paradoxical conclusion — and a comparative form that is used as the basis for the inductive step. In fact, vague predicates — and in particular GRADABLE ADJECTIVES like *tall, big, fast* and *heavy* — generally have more than just positive and comparative forms; they appear in a host of complex constructions that are used to express different "degrees" to which the predicate applies to its argument, some of which are illustrated in (2).

(2) a. Kim is very/really/rather/quite tall.
b. Kim is six feet tall.
c. Kim is less tall than Lee.
d. Kim is tall compared to Lee.
e. Kim is as tall as we expected her to be.
f. Kim is too tall to fit in this seat.
g. Kim is tall enough to reach the ceiling.
h. Kim is so tall that she has to buy special clothes.

The challenge presented by these examples for semantic theory is to come up with an analysis of the meanings of words and phrases like *tall*, *heap of sand*, *likes broccoli* and *near the ocean* that simultaneously explains their semantic and pragmatic properties in the unmarked, positive form and provides the basis for a compositional account of comparatives and the other complex constructions in (2). Central to this challenge is an explanation of the fact — already evident from a close examination of (1) — that the positive form is vague (in the sense of being incompatible with sharp distinctions) but the comparative form is not, at least for monodimensional predicates like *tall(er)*.[1] In setting up the scenario in which we discover that we are unwilling to judge some x and y differently for α when they are minimally distinct with respect to α — that α is necessarily tolerant, to use Wright's (1975) terminology — we crucially rely on the fact that we are willing to judge them differently for *more* α *than*.

The story of the Sorites Paradox in linguistics — and the story of vagueness in linguistics more generally — is intertwined with the story of the semantic analysis of constructions like those in (2), in particular comparatives, and of the analysis of the grammar of gradability more generally, especially the grammar of gradable adjectives. This work can be (somewhat coarsely) divided into two traditions which differ, essentially, in whether vagueness or comparison is taken to be basic. In the first approach, which emerges out of work on the Sorites Paradox in philosophy of language, vagueness is basic, and a semantics for gradability is built on top of a logic for vagueness. In the second approach, which is more rooted in syntax and linguistic analysis, gradability is taken to be basic, and vagueness is derived. In sections 1.2 and 1.3, I describe these analytical traditions in more detail, with close attention to the empirical

[1] It is less clear that the same distinction holds between multidimensional predicates like *clever* and *good*. To keep the discussion focused, I will limit my attention to monodimensional gradable predicates for most of this chapter, but I will say a bit more about multidimensional gradable predicates in section 1.5.

arguments that both motivate and challenge them, and to the ways that they inform, and are informed by, theorizing about the Sorites Paradox. Section 1.4 discusses a second class of gradable adjectives that present a challenge for existing analyses of the relation between gradability and vagueness, because they have positive forms that do not give rise to the Sorites Paradox, and section 1.5 concludes with a discussion of recent work in linguistics that is specifically geared towards understanding the Sorites in terms of communication in the presence of semantic uncertainty.

1.2 From vagueness to comparison

The first attempts within linguistics to provide a compositional analysis of the semantics of comparative constructions built directly on work in the philosophical literature geared towards understanding vagueness and the Sorites Paradox. These attempts are similar in treating comparatives as expressions that manipulate or constrain parameters relevant for fixing the extension of the positive form, but differ in the details of what this parameter is.

The first kind of analysis is briefly gestured at by Lewis (1970) and then worked out in substantive detail, alongside a comprehensive syntactic analysis of comparatives, by McConnell-Ginet (1973). Following Lewis, McConnell-Ginet assumes that the extension of a vague predicate is determined by an element of the index of evaluation (an array of parameters relative to which extensions are fixed) called a DELINEATION COORDINATE, which is a value in an ordering appropriate to the type of concept the predicate describes (height, weight, temperature, etc.) that represents the boundary between the things the predicate is true and false of. The denotation of *tall* on this view, for example, is as stated in (3). (Here and throughout, $[\![\alpha]\!]^{\pi}$ means "the extension of α (an expression of the object language) relative to parameter π.")

(3) $[\![tall]\!]^{d} = \{x \mid x's \text{ height is at least as great as } d\}$

The comparative can then be analysed as an expression that quantifies over delineations: *taller than* expresses a relation between individuals such that there is a delineation that makes *tall* true of the first and false of the second.[2] More generally, for any gradable adjective α:

[2] Note that this analysis says nothing about whether a delineation is or is not

(4) $[\![more\ \alpha]\!]^d = \{\langle x, y\rangle \mid \exists d'[[\![\alpha]\!]^{d'}(x) \wedge \neg[\![\alpha]\!]^{d'}(y)]\}$

Kamp (1975) develops an analysis of comparatives that is conceptually identical to the Lewis/McConnell-Ginet analysis, but is formalized in terms of a supervaluationist approach to vagueness (cf. Fine, 1975). (The key difference between the two approaches is that the latter is partial, allowing for extension gaps, but the former is not.) In Kamp's analysis, rather than quantifying over delineations, comparatives quantify over the (possibly partial) models relative to which vague predicates are assigned extensions. Given a general "consistency constraint" which ensures that for any objects x, y, if x is in the positive extension of α relative to some model M and y is not (see Klein 1980; van Benthem 1982), there is no M' such that y is in the extension of α and x is not, the semantics of the comparative can be stated as in (5).

(5) $[\![more\ \alpha]\!]^M = \{\langle x, y\rangle \mid \{M' \mid [\![\alpha]\!]^{M'}(x)\} \supset \{M' \mid [\![\alpha]\!]^{M'}(y)\}\}$

Taller than, on this view, expresses a relation between individuals x, y in a model M such that the models that make *tall* true of x are a proper superset of the models that make *tall* true of y. Note that it need not be the case that *tall* is true of x or y in M: *tall* could be false or undefined for one or both of them (the latter for borderline cases). Given the consistency constraint, the truth conditions of *taller than* effectively require just that there be some model in which *tall* is true of x and false of y, similar to what we saw with the Lewis/McConnell-Ginet analysis.

A hybrid of these two approaches is developed by Klein (1980) (see also Wheeler 1972; van Benthem 1982; van Rooij 2011a,b; Burnett 2016), who analyzes gradable predicates as expressions whose extensions are determined relative to subsets of the domain of discourse called COMPARISON CLASSES. Comparison classes provide the domain of the predicate and are fixed by the context, like the delineation coordinate. They also provide the basis for fixing the extension of the predicate, as a function of the way that objects in the comparison class distribute relative to the grandable concept the predicate is used to describe (height, weight, etc.), in a way that allows for partiality, like the partial models in a supervaluationist analysis. Finally, interpretations of gradable predicates are also subject to a variant of the consistency constraint, such that for any

usable in the positive form. For example, it may be the case that that there is no context in which the two smallest things in the universe would ever be appropriately described as tall. But the analysis is committed to the position that if they differ in height, there is a delineation that makes *tall* true of one of them and false of the other.

x, y, if x is in the positive extension of α relative to some comparison class C and y is not, there is no C' such that y is in the extension of α and x is not, and also to a general informativity principle that stipulations that for any C, the positive and negative extensions of α relative to C must both be non-empty. Since the comparison class provides the domain of the predicate, this latter constraint rules out singleton comparison classes, and together with consistency allows for the two ways of defining the semantics of comparatives shown in (6a-b), due to Klein (1980) and Wheeler (1972) respectively.

(6) a. $[\![more\ \alpha]\!]^C = \{\langle x, y\rangle \mid \exists C'[[\![\alpha]\!]^{C'}(x) \wedge \neg[\![\alpha]\!]^{C'}(y)]\}$
b. $[\![more\ \alpha]\!]^C = \{\langle x, y\rangle \mid [\![\alpha]\!]^{\{x,y\}}(x) \wedge \neg[\![\alpha]\!]^{\{x,y\}}(y)\}$

These different approaches to comparatives are similar in that each of them derives the meaning of the complex form *more* α in terms of the meaning of α, and in particular each makes crucial use of the theoretical machinery brought to bear in the analysis of the vagueness of α to do so. Kamp's approach is most directly tied to a particular account of the Sorites Paradox, since it is committed to a supervaluationst semantics; the other two accounts are compatible with supervaluationism but are equally compatible with e.g. contextualist or epistemicist approaches to the Sorites.

These analyses are also similar in their responses to the challenge of explaining why the comparative form is not vague, even though the positive form from which it is derived is vague. In each account, vagueness in the positive form is ultimately tied to indeterminacy or uncertainty about the parameter relative to which its extension is determined in a context of utterance: the delineation coordinate, the model, or the comparison class. In the case of comparatives, however, these parameters are fixed in a fully determinate and compositional way.

Note, however, that it follows from the fact that the meaning of the comparative is stated in terms of the meaning of the positive that each of these analyses must be committed to the position that the positive form, contrary to normal appearance and use, is in fact compatible with a precise meaning (as noted by both Kamp 1975 and Klein 1980). To see why, consider the different interpretations assigned by these analyses to the sentence in (7), which are spelled out in (8).

(7) Kim is taller than Lee.

(8) a. $\exists d'[[\![tall]\!]^{d'}(\mathbf{k}) \wedge \neg[\![tall]\!]^{d'}(\mathbf{l})]$ McConnell-Ginet
b. $\{M' \mid [\![tall]\!]^{M'}(\mathbf{k})\} \supset \{M' \mid [\![tall]\!]^{M'}(\mathbf{l})\}$ Kamp

c. $\exists C'[[\![tall]\!]^{C'}(\mathbf{k}) \wedge \neg[\![tall]\!]^{C'}(\mathbf{l})]$ Klein

d. $[\![tall]\!]^{\{\mathbf{k},\mathbf{l}\}}(\mathbf{k}) \wedge \neg[\![tall]\!]^{\{\mathbf{k},\mathbf{l}\}}(\mathbf{l})$ Wheeler

In a context in which Kim's height is greater than Lee's height by a very small amount — say 5 millimeters — (7) is both true and felicitous, regardless of how tall either of them actually is, which in turn means that there must be delineations/precisifications/comparison classes relative to which Kim is in the extension of *tall* and Lee is not. It follows that on all of these views, the tolerance of the positive form cannot be "hardwired" into the lexical semantics of gradable predicates, since the positive form is, by hypothesis, the lexical entry.

But this in turn leads to apparently paradoxical results when we look beyond "explicit" comparison constructions involving morphologically marked adjectives, such as (7), to instances of "implicit" comparison constructions involving the positive form, such as (9) (Sapir, 1944; Kennedy, 2007a, 2011; van Rooij, 2011a).

(9) Kim is tall compared to Lee.

(9) entails that Kim's height exceeds Lee's height, just like (7), a result that is most plausibly obtained by assuming that the semantic function of *compared to* is to cause the denotation of *tall* to be fixed in such a way as to make it true of Kim and false of Lee, just like the comparative. But (9) differs from (7) in that it is infelicitous in the context described above, in which Kim and Lee differ in height by a very small amount, and indeed in any context which reproduce the "minimal distinctions" property of soritical reasoning. ((9) is also different from (7) in implicating that Kim is not tall, as pointed out by Sawada 2009.)

Positive adjectives used as modifiers in definite descriptions show similar behavior. In a context in which I am asked to identify two individuals standing next to each other in front of a wall which indicates that their heights are 1.75 meters and 1.5 meters respectively (like a police lineup), either of the utterances in (10) would be felicitous ways for me to say who is who.

(10) a. The tall one is Kim; the short one is Lee.

b. The taller one is Kim; the shorter one is Lee.

Note that (10a) is acceptable regardless of whether a height of 1.75 meters would otherwise be sufficient to justify characterizing Kim as tall (e.g. if the comparison class is adult American males). This is because, as we have seen, positive *tall* is context dependent, and in a context involving just two individuals, its meaning can be fixed in a way that makes it true of one individual and false of the other — in this case as a way of satisfying the uniqueness presuppositions of the definite determiner. However, if the context is one in which the heights of the two individuals are very similar — e.g., if one is 1.75 meters tall and the other is 1.745 meters tall, as in the scenario described above — only the utterance involving the morphologically comparative form in (10b) is acceptable.

In sum, the problem of implicit comparison is that the very same assumptions which enable a compositional analysis of the comparative form of an adjective in terms of the meaning of the positive form lead to a paradox. The truth and felicity of (7) and (10b) in the minimal distinctions context means that there must be delineations, precisifications or comparison classes that make *tall* true of Kim and false of *Lee*; yet the infelicity of (9) and (10a) means that these delineations, precisifications and comparison classes cannot be used to fix the extension of *tall* in a minimal distinctions context. Evidently, the relevant valuations are compositionally manipulable by the comparative morphology, but otherwise inaccessible to the positive form, even when other features of the utterance, such as the meaning of *compared to* or the presuppositions of the definite article, should promote their use. The challenge for proponents of the semantic analyses of positive and comparative adjectives described in this section is to derive these results in a principled way. (See van Rooij 2011a for one attempt to do so, and Kennedy 2011 for a response.)

The explicit/implicit distinction also helps further define the evaluation criteria for accounts of the Sorites Paradox. If it is assumed that one of the analyses outlined here (or the equivalent) correctly characterizes the semantics of vague predicates, then tolerance — whatever is responsible for judgments about the inductive premise — cannot be a matter of lexical semantics, since the lexical semantics of vague predicates (or at least positive form gradable adjectives like *tall* and *bald*) is compatible with precise interpretations. It must instead have something to do with constraints on how a vague predicate can be interpreted or justifiably used in contexts of utterance, or with properties of contexts, or perhaps with the interaction between contexts and the semantics of generalizing statements like the inductive premise — as in e.g. epistemicist and contextualist accounts of vagueness. But whatever the account is, it should also say something about why tolerance obtains not only in generalizing statements like the inductive premise, but also in implicit comparison constructions, in which we might otherwise expect the semantic and pragmatic properties of the rest of utterance ought to make such "inaccessible" contexts or interpretations accessible.

1.3 From comparison to vagueness

The analyses of comparatives described in the previous section all share the feature of defining the meanings of complex degree constructions in terms of the meanings of the positive form, and in particular in terms of the analytical assumptions that are involved in accounting for the vagueness of the positive form. In this sense, these analyses are directly connected to work in philosophy of language on the Sorites Paradox. Alongside these analyses we find a second set of approaches that are not connected to work on vagueness or the Sorites, but are instead more closely connected to work in semantic theory on the compositional analysis of other linguistic phenomena. This line of work effectively takes the complex constructions illustrated in (2) as the starting point, and asks what kind of lexical meaning for a bare gradable adjective like *tall* best captures the full range of meanings of the complex forms, without any special concern for the vagueness of the positive form or the proper account of the Sorites Paradox.

The answer that this line of work provides is that the various complex forms are not derived directly from the positive form, but instead both the positive and comparative (and other complex) forms are derived from a more basic lexeme that denotes a relation between individuals and some other value — typically some kind of measurement value, or DEGREE — as in (11). (See e.g. Bartsch & Vennemann 1972; Seuren 1973; Cresswell 1976; von Stechow 1984; see Klein 1991 and Bale 2009 for detailed discussions of how degrees can be modeled in terms of equivalence classes of individuals, rather than as independent model-theoretic objects.)

(11) $[\![tall]\!] = \{\langle x, d\rangle \mid x's \text{ height is at least as great as } d\}$

On this view, to be a gradable predicate is just to denote a relation between individuals and degrees. The function of the complex constituents that combine with the adjective in examples like (2) is to fix the value of the degree argument, thereby providing a standard of comparison, and turning the relation into a property that holds of an object just in case it has a degree of the relevant property that is at least as great as that standard. The predicate in (2b), for example, denotes the property of having a height that is at least as great as the degree denoted by *six feet*; the comparative of inferiority in (2c) denotes the property of having a height that is exceeded the degree of Lee's height; the comparative of superiority simply switches the ordering relation; and so on.

The similarity between (11) and the delination-based semantics for gradable predicates in (3) is not accidental. In effect, degree-based (and other relational) analyses take the lexical meanings of gradable predicates to be de-contextualized variants of the meanings assumed by the vagueness-based approaches discussed in the previous section, where the interpretive parameter used to fix the extension of a vague predicate is reanalyzed as one of the arguments of the relation. The fact that comparatives are not vague, on this view, is unremarkable, since their meanings are derived from forms with fixed meaning like (11), which are also not vague. In short, whether a predicate built out of a gradable adjective is vague or not does not depend on the denotation of the adjective, but rather on the meaning introduced by the degree morphology with which it composes.

For this very reason, it is the vague positive form that, somewhat paradoxically, presents the trickiest analytical challenge, for two reasons. The first is a compositional one. The denotation of *tall* in (11) expects two arguments: a degree-denoting one and an individual-denoting one. But there is no constituent in the surface form of a sentence containing positive *tall* to saturate the degree argument and thereby convert it from a relation between degrees and individuals to a property of individuals, which is the semantic type it must have when it is used as a modifier or predicate. The usual response to this problem is to hypothesize a phonologically null, "positive degree" morpheme which does this job, or to assume default saturation by a free variable over degrees, or default existential binding of the degree argument. These are all analytical options that are independently motivated for other cases of "implicit argument" saturation, but appeal to such operations in one language are typically motivated by pointing to another language in which the hypothesized argument slot is obligatorily realized by overt morphosyntax. The same cannot be said for the positive form of a gradable adjective: there is no language in the world that we know of in which a gradable adjective must compose with some overt morphosyntax to create a predicate with positive form meaning (Klein, 1980; Francez & Koontz-Garboden, 2015; Grano & Davis, to appear). This is deeply mysterious if gradable adjectives denote relations rather than properties of individuals — if they are basically two-place rather than one-place — and it is a mystery that so far has not been satisfactorily resolved by proponents of relational analyses of gradable adjectives.

Alongside the compositional challenge is the question of how exactly to characterize the semantics of the positive form, once a particular assumption about how composition actually works has been adopted, and the further question of whether the semantic analysis adopted provides insights on vagueness or the Sorites Paradox. The latter question is particular salient in the relational analysis, given the fact that the lexical meaning of a gradable predicate such as (11) is not itself vague, but it is a question that, until fairly recently, has not received a great deal of attention from scholars in linguistics, for understandable reasons. As noted above, the relational approach emerged out of a research program aimed at understanding and giving compositional analyses of complex degree constructions, which manifest a great deal of grammatical complexity and variation both within and across languages (see e.g. Stassen 1985; Beck et al. 2009; Bochnak 2013). The semantics of the grammatically simple positive form is usually presented as an afterthought — if it is addressed at all — and discussion of the implications of a particular semantic analysis for vagueness and the Sorites Paradox is rare.

The simplest analytical option is simply to assume that composition of the relational meaning in (11) with a null positive morpheme reproduces one of the lexical semantic analyses described above, such that the degree argument of the adjective is constrained to exceed a threshold degree whose value is fixed by a contextual parameter, as in delineation analyses (see e.g. Barker, 2002; Kennedy & McNally, 2005), possibly in a partial way, as in supervaluationist analyses (see e.g. Sassoon, 2009, 2013).[3] These analyses do not provide new options for analyzing the Sorites Paradox that are not already present in corresponding approaches that treat the positive form as basic, but they do provide the basis for a more explanatory account of implicit comparison, if not a fully satisfactory one.

Recall from above that in a context in which Kim and Lee differ minimally in height, (12a) is infelicitous but (12b) is fine.

(12) CONTEXT: Kim is 90cm tall; Lee is 89.5cm tall.
  a. # Kim is tall compared to Lee.
  b. Kim is taller than Lee.

As we have already seen, if the comparative is derived from the positive,

[3] Actually, neither Barker nor Sassoon introduce degrees into the object language, but both make use of degrees in the model theory in ways that reproduce the relational analysis in the metalanguage.

this difference between explicit and implicit comparison is puzzling. The infelicity of (12a) indicates that there is no way of fixing the meaning of *tall* that makes it true of Kim and false of Lee when they differ minimally in height, as would expected if *tall* were necessarily tolerant. But the felicity of (12b) requires that there be such a way of fixing the meaning of *tall* in such a context, since this is what is required by the meaning of the comparative, which means that *tall* is not necessarily tolerant.

If, however, the comparative is not derived from the positive, and instead both positive and comparative are separately derived from a more basic relational meaning like (11), this paradox disappears. The meaning of the adjectival predicate in (12a) is a function of (11) and a phonologically null positive morpheme, so we may assume from the infelicity of (12a) that, whatever the semantic contribution of the positive morphology is, the result of composition with (11) is indeed a meaning that is necessarily tolerant — regardless of whether it is used in an implicit comparison construction like (12a) or in the inductive premise of the Sorites Paradox. But this assumption makes no additional predictions about (12b), since the meaning of the adjectival predicate in this example is a function of (11) and the comparative morpheme. As long as the comparative morpheme derives a meaning that is not subject to tolerance — e.g., if it simply entails that the degree to which Kim is tall exceeds the degree to which Lee is tall — we can accommodate the difference between (12a) and (12b). Of course, these assumptions do not yet provide an explanation of why the positive is subject to tolerance, but they draw a compositional distinction between implicit and explicit comparisons that cannot be drawn in non-relational analyses. Moreover, our compositional semantic assumptions do not commit us to a particular analysis of tolerance: since we need not assume that the positive form ever has a precise interpretation, we open the door to an analysis of tolerance that is rooted in its semantics — or more precisely, to the semantics of the positive morpheme. Specifically, since the positive form is compositionally derived in a relational analysis, it is possible to introduce explicit semantic content into the truth conditions of the positive form that is not shared by comparative and other degree-modified forms, via the semantic content of the positive morphology; this is an option that is unavailable in lexicalist analyses, though as we will see, not all compositional analyses take advantage of it.

For example, one common approach along these lines, geared towards encoding the intuition that the positive is true only of objects that fall above some threshold in a distribution, assigns a denotation to the positive morpheme which restricts the degree argument of the adjective to degrees that exceed an average, norm or some other value that is a function of the distribution of objects in an explicit or implicit comparison class on an adjective-appropriate scale (see Bartsch & Vennemann 1972, 1973; von Stechow 1984; Bierwisch 1989 and Kennedy 1999 for variants of this idea). Another strategy is to analyze the positive morpheme in such a way that it constrains the degree argument of the adjective to exceed a contextually-determined "neutral zone" on a scale (see e.g. Heim, 2006; von Stechow, 2009). When combined with a theory of antonymy that characterizes pairs like *tall* and *short* as encoding inverse ordering relations, this kind of analysis semantic encodes the intuition that there are regions on e.g. the height scale in which objects count as neither tall nor short.

Like their "content light" cousins which merely reproduce lexicalist denotations for the positive via semantic composition, these "content heavy" analyses of the positive form do not provide new options for analyzing the Sorites Paradox, and arguably fare worse than the former in helping us understand the appeal of the inductive premise — and answering what Fara (2000) calls the "psychological question" about the Sorites — since the denotations they provide for the positive form are less indeterminate than denotations that reproduce lexicalist analyses. The truth conditions of (13a) on a "greater than average" semantics for the positive, for example, are as in (13b). But the latter is easily rejected, while the former is not.

(13) a. If x is tall and there is a y such that x is just a bit taller than y, then y is tall.

b. If x's height is at least as great as the average height for class X, and there is a y such that x's height is just a little bit greater than y's height, then y's height is at least as great as the average height for class X.

Likewise, these analyses have a hard time accounting for constraints on implicit comparison. A very natural way of thinking of the function of the *compared to* phrase in (12a), for example, is that it restricts the comparison class to the set consisting of just Kim and Lee. But if that is the case, then even if Kim's height is only slightly greater than Lee's it will still be the case that it exceeds the average height for the class {Kim, Lee}.

The problem with these approaches, in a nutshell, is that they make the semantic content of the positive too similar to that of the comparative, and so too determinate, leaving very little room for vagueness

(Bogusławski, 1975). The alternative is to assume that the positive introduces semantic content that results in vague truth conditions. Cresswell (1976), for example, characterizes the positive as requiring the degree argument of the adjective to be "towards the top of the scale;" Fara (2000) argues that it should "significantly exceed an average" (cf. Solt, 2009); Bogusławski (1975) says it should be "noteworthy;" and Kennedy (2007b) says that it should represent a degree that "stands out" relative to the kind of scale the adjective uses. Of all the semantic analyses of the positive within the relational tradition, these approaches are most sensitive to the Sorites Paradox, and to providing an account of tolerance more generally. Fara (2000), for example, argues that the "significantly greater than" relation is interest relative in a way that derives tolerance, and Kennedy (2007b) provides a pragmatic account of why it can never be the case that for two objects that differ minimally along a scalar dimension, one object can be judged to have a degree of the relevant property that "stands out" while the other does not (with two key exceptions, that I will return to below). In these kinds of analyses, the explanation of intuitions about the inductive premise of the Sorites and the explanation of the infelicity of implicit comparisons in minimal differences contexts are the same: both emerge from the semantic content of the positive form, which gives rise to tolerance.

However, these kinds of accounts face a challenge of undergeneration when we look beyond the core cases of soritical adjectives to a wider range of data. As I pointed out in section 1.1, soritical arguments can be constructed using expressions from a range of different grammatical categories, not just from gradable adjectives. An account of tolerance that is based on the details of the compositional analysis of gradable adjectives, via the semantics of the positive form morphology, says nothing about these other cases.[4] Such an analysis is bound to be incomplete as a full explanation of vagueness, then, no matter how well it explains the particular case of gradable adjectives.

That said, an analysis in which tolerance emerges from the semantics of the positive form is of course consistent with the possibility that it can also arise in other ways, from other features of meaning or use. The argument for the relevance of semantic content is rooted in linguistic facts,

[4] With a couple of interesting exceptions (such as certain kinds of verbs describing changes of state; see Kennedy & Levin 2008; Bobaljik 2012), there is no empirical evidence in support of generalizing the (de-)compositional analysis of the positive form of gradable adjectives to all vague expressions, regardless of grammatical categories.

namely the difference in acceptability between explicit and implicit comparison constructions in minimal difference contexts, which is difficult to explain on accounts that do not tie tolerance to the content of the positive form. If this argument is correct, then the fact that tolerance is also observed in constructions that do not involve positive form gradable adjectives means that the origins of tolerance are heterogenous, and our best theory will be one that can explain what unifies the different semantic, pragmatic or other factors that bring it about.

1.4 Absolute adjectives

The discussion so far has focused on different attempts to capture a key linguistic property of vague predicates that is evident in the Sorites Paradox: the fact that such predicates are obligatorily tolerant, but have compositionally derived comparative forms which are not. And as we have seen, in the case of gradable adjectives, the analysis of comparison — and gradability more generally — is closely tied to the analysis of tolerance. Approaches such as the ones discussed in section 1.2 take the positive form of a gradable adjective to be basic, assume an essentially tolerant meaning for it, and derive gradability from this meaning (in terms of the partitioning of the domain relative to different valuations of the predicate allowed by the consistency constraints). Approaches such as those discussed in section 1.3, on the other hand, assume an essentially gradable, relational meaning as the lexical content of the adjective, and derive the tolerance of the positive form compositionally, in the mapping of the underlying relational meaning to a property. In both types of analysis, then, the expectation is that if a predicate is gradable — if it can appear in comparatives and other degree constructions like the ones illustrated in (2) — then it should display tolerance in its positive form, and should give rise to the Sorites Paradox.

As it turns out, there is a class of gradable adjectives that, like other members of this class combine readily with comparative and other degree morphology, but which need not display the kind of tolerance in their positive forms that all of the analyses discussed so far would lead us to expect (Rusiecki, 1985; Cruse, 1986; Rotstein & Winter, 2004; Kennedy & McNally, 2005; Kennedy, 2007b). Adopting terminology from Unger 1975, Kennedy & McNally (2005) refer to such adjectives as ABSOLUTE gradable adjectives, and distinguish two kinds. Maximum standard absolute adjectives like *straight, empty, dry* and *flat* require their arguments to have a maximal degree of the relevant property, and minimum standard absolute adjectives like *bent, wet, scratched* and *bumpy* merely require their arguments to have some degree of the relevant property.

These properties of absolute adjectives are evident in entailment relations between comparatives and positives (Kennedy, 2007b). As shown by (14), ordinary gradable adjectives like *tall* and *short* — which are also referred to as RELATIVE gradable adjectives when it is important to distinguish them from their absolute cousins — lack entailments from comparative to positive:

(14) a. x is taller than $y \not\Rightarrow x/y$ is (not) tall
b. x is shorter than $y \not\Rightarrow x/y$ is (not) short

Absolute adjectives, however, display a different pattern. A comparative involving a maximum standard adjective like *straight* entails the negation of the positive for its second argument, and a comparative involving a minimum standard adjective like *bent* entails the corresponding positive for its first argument:

(15) a. x is straighter than $y \Rightarrow y$ is not straight
b. x is more bent than $y \Rightarrow x$ is bent

The pattern in (14) is what we expect on any of the theories of the positive form discussed in sections 1.2 and 1.3: the comparative effectively requires its arguments to differ in height, such that the first has more/less than the second, but says nothing about whether their heights are significant, stand out, exceed an average, or anything else (relative to the different orderings that the antonyms impose). The pattern in (15), on the other hand, is not expected on any of the views we have discussed so far, and instead indicates that positive *straight* requires its argument to be maximally straight (and so is false of y in (15a), since y must be less straight than x), and that positive *bent* requires its argument to merely have some degree of bend (and so is true of x, which has some bend if it is more bent than y).

This has consequences for the Sorites Paradox: as pointed out by Kennedy (2007b), absolute gradable adjectives do not lead to paradox in the same way as relative gradable adjectives. More precisely, although we might be inclined to accept (16a) and (16b) in most contexts, we may also reject these claims in a way that is impossible in soritical arguments involving relative adjectives like *tall* and *short*.

(16) a. For any x, if x is straight, and there is a y such that x is just a bit straighter than y, then y is also straight.

b. For any x, if x is bent, and there is a y such that x is just a bit more bent than y, then y is also bent.

Put another way, while relative adjectives are inherently tolerant, absolute adjectives allow for what Pinkal (1995) calls NATURAL PRECISIFICATIONS: (16a) can be rejected because *straight* can be used to mean (something like) "perfectly straight," and (16b) can be rejected because *bent* can be used to mean (something like) "minimally bent."[5] This is also shown by the contrast between absolute and relative adjectives in implicit comparison contexts. The two lines in Figure 1.1 differ minimally along two dimensions: degree of bend and degree of length. (17a) is infelicitious, as we have seen, because *long* and *short* cannot be used to distinguish between objects that differ minimally in length (as is required, in this case, by the presuppositions of the definite description). In contrast, (17b) is felicitous because *straight* and *bent* can be used to distinguish between two objects that differ minimally in degree of bend.

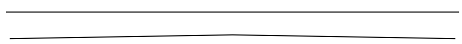

Figure 1.1

(17) a. # The long line is straighter than the short line.
b. The straight line is longer than the bent line.

If facts like those above are interpreted as indicating that the lexical meanings of absolute adjectives are essentially fixed, then the challenge they present for approaches that derive gradability from a semantics for tolerance is to explain where gradability comes from. The most articulated response to this challenge is provided by Burnett (2016), who bites the bullet and assigns absolute adjectives precise meanings — *straight*, for example, denotes the property of being perfectly straight — but provides a means of deriving tolerant meanings pragmatically in a way that can then support comparative formation.

[5] Philosophers might object at this point that nothing (with the relevant physical properties) is perfectly straight, and correspondingly that everything (again with the relevant properties) has minimal bend, rendering such meanings unusuable. And indeed, actual uses of *straight* and *bent* communicate something more like "relatively close to straight" or "relatively bent." But this makes it all the more remarkable that absolute adjectives can be used to make sharp distinctions in ways that relative adjectives cannot. Whether "imprecise" uses of absolute adjectives are a function of semantic content or pragmatic reasoning is a question of current debate, relevant not just for gradable adjectives but also numerals, measure terms, temporal phrases, place names, and so forth. See Lasersohn 1999; Krifka 2007; Syrett et al. 2010; Lassiter & Goodman 2014; Burnett 2016; Leffel et al. to appear; Klecha 2017 for discussion and analyses.

The challenge of absolute adjectives for approaches that derive tolerance from gradability is to explain why their positive forms have precise interpretations based on maximal or minimal thresholds rather than the tolerant, "norm-based" denotations that are assigned to positive form relative adjectives. Responses to this challenge have made key use of a lexical semantic difference between classes of adjectives that relational analyses are particularly well-equipped to model: gradable adjectives differ in whether they encode scalar concepts that are based on open or closed scales. This distinction can be diagnosed by looking at acceptability with certain types of modifiers (Rotstein & Winter, 2004; Kennedy & McNally, 2005; Kennedy, 2007b; Syrett, 2007). The modifier *completely*, for example, introduces the entailment that an object has a maximal degree of a gradable property, and so combines only with adjectives that use scales with maximum values, while the adjective *slightly* entails that an object exceeds a minimum degree, and so selects for adjectives that use scales with minimum values. As the following examples show, there is a correlation between the relative/absolute distinction and scale structure: absolute adjectives have closed scales; relative adjectives have open scales.[6]

(18) a. completely straight/empty/flat — ABSOLUTE MAX
b. # completely long/heavy/big — RELATIVE

(19) a. slightly bent/open/striped — ABSOLUTE MIN
b. # slightly long/heavy/big — RELATIVE

Two different kinds of explanations for this correlation have been proposed in the literature. In the first approach, scale structure plays a direct role in the determination of thresholds of application and the relative/absolute distinction. The maximum and minimum degrees on closed scales provide salient values for coordination on thresholds, and threshold interpretation is optimized to select such values whenever possible (Kennedy, 2007b; Potts, 2008; Toledo & Sassoon, 2011; Qing & Franke, 2014a,b). Open scales, in contrast, lack any "natural transitions" (to use Williamson's (1992) term) that could provide coordination points

[6] The examples in 19b are crucially unacceptable on interpretations that are parallel to the most prominent interpretations of the examples in 19a, which would be paraphrased as "a slight amount of length/weight/size." These examples can have a different kind of interpretation, paraphrasable as "slightly *too* long/heavy/big," i.e. as expressions of slight excess. But in such cases the semantics of excess provides a minimum standard for the modifier to interact with, namely the minimum degree that counts as excessive for the relevant purpose.

for thresholds, and so have highly context dependent and uncertain interpretations. In the second kind of approach, the role of scale structure is more indirect. Decisions about thresholds are based not on formal semantic properties of gradable adjectives, but rather emerge from Bayesian reasoning that takes into account probabilisitic prior knowledge about the distribution of objects in a gradable adjective's domain relative to its scale (lg14,, lg15; the basic mechanisms of this approach are described in more detail in the next section). The domains of open scale adjectives tend to have normal distributions (our experience with nails, for example, tells us that there are relatively few very short and very long ones, but lots in between), which give rise to similarly shaped curves for threshold probabilities and relative interpretations. The domains of closed scale adjectives have non-trivial probability mass at the scalar endpoints (there are lots of straight or nearly straight nails), which give rise to threshold probabilities that are much narrower and skewed towards the endpoints, giving rise to absolute interpretations — albeit ones that are "less absolute" than in the first kind of approach.

1.5 Recent developments

The two approaches to the relative/absolute distinction discussed in the previous section differ in the role that scalar distinctions play in determining meaning, but they share the core intuition that this distinction — which is, at its core, the question of whether a predicate is necessarily tolerant or not — emerges from the mechanisms involved in determining the communicative content of utterances that introduce semantic uncertainty. In the case of positive form gradable adjectives, the uncertainty is about the threshold that an object must reach in order to count as having the property in question. Absolute interpretations show that the principles involved in fixing thresholds are optimized to minimize uncertainty (based on scalar representation, domain distribution, or perhaps a combination of both); relative interpretations, and tolerance, result when threshold uncertainty cannot be minimized.

Threshold uncertainty clearly does not present a problem to communication, however, and one of the important contributions of Lassiter and Goodman's (2014; 2015) work is to provide a model for how communication with vague predicates actually works, and in so doing provide new insights on the Sorites Paradox. The core idea of Lassiter and Goodman's model is that a "pragmatic listener" computes a probability distribution for the value of an uncertain variable in an utterance as a function of a relevant set of prior values and the assumption that speakers choose particular utterances (over relevant alternatives) with (at least) a goal of optimizing informativity. In the case of gradable adjective interpretations, the model uses prior knowledge of the distribution of the degrees to which the objects in the adjective's comparison class possess the scalar concept associated with the adjective to derive a probability distribution on thresholds, and a corresponding posterior probability about the scale position of an object that is described as having the property expressed by the adjective.

As illustration, suppose that (20) is uttered in a context in which the comparison class consists of the sorts of nails that are encountered in typical household carpentry contexts, the lengths of which are approximately normally distributed.

(20) That nail is long.

The further below the mean that a particular length is, the more likely it is that an arbitrary nail has at least that much length, and the further above the mean that a particular length is, the less likely it is that an arbitrary nail has that length. Given the speaker's preference for informativity, a low value for the adjective's threshold of application (e.g. one that makes *long* true of 75% of the comparison class) will be assigned low probability, because the resulting meaning would be too weak, while a high value for the threshold (e.g. one that makes *long* true of only 1% of the comparison class) will also be assigned a low probability, because the resulting meaning would be too strong. The output of the model in a simple case like this is a posterior probability distribution over thresholds that is shifted upwards from the prior distribution over the comparison class, and a posterior probability for the length of the target of predication that is shifted still further up the scale.

Put another way, the model predicts that (20) will be heard to communicate something roughly equivalent to "the length of this nail is significantly greater than the average length of nails in the comparison class," which, as we have seen, is an accurate paraphrase of its truth conditions. But more than that, Lassiter & Goodman (2015) show that the model also makes specific quantitative predictions about the plausibility of the inductive premise of the Sorites Paradox, and the extent to which it should be heard as compelling, which vary as a function of different semantic parameters, such as the granularity of the scale and different ways of modeling the semantics of the inductive premise of the

Sorites (in particular involving different approaches to the semantics of the conditional). This work not only provides a fully explicit account of communication with, and intuitions about, vague language; it also opens up the possibility of subjecting different accounts of the Sorites to experimental investigation.

Lassiter and Goodman's model is rooted in the idea that tolerance reflects uncertainty about threshold values; recent work by Grinsell (2017) develops the hypothesis that tolerance emerges from uncertainty about the factors that are involved in making decisions about thresholds in the first place. Focusing on MULTIDIMENSIONAL adjectives like *healthy* and *talented*, Grinsell observes that such expressions have meanings that are formally comparable to multidimensional decision problems of the sort studied by Arrow (1950, 1959) in his work on social choice. Just as the decisions of a legislature involve an aggregation of the individual choices of multiple agents with distinct and possibly conflicting preferences, the extension of a predicate like *healthy* involves an aggregation of choices along multiple, independent dimensions of measurement, such as blood pressure, body temperature, weight, and so forth. Arrow proved that given a certain reasonable set of constraints on the decision making process, the aggregation of set of individual choices into a single multidimensional choice function is guaranteed to fail transitivity — it will allow for rankings of A over B, B over C, and C over A — even if the rankings produced by the individual dimensions are linear. Grinsell argues that the same set of Arrowian principles constrain the denotations of multidimensional predicates, with the same result. For example, if we know that Kim has better blood pressure than Lee and Lee is less overweight than Pat, we may judge (21a) and (21b) to be true, but it does not guarantee that (21c) is true.

(21) a. Kim is healthier than Lee.
b. Lee is healthier than Pat.
c. Kim is healthier than Pat.

According to Grinsell, the Sorites Paradox emerges as a concrete manifestation of the Arrowian constraints on multidimensional meanings. On the one hand, these constraints entail that the tolerance relation expressed in the comparative form of a multidimensional predicate is not transitive, in which case the inductive premise of the Sorites is not guaranteed to be true, and the paradoxical conclusion does not follow. At the same time, the Arrowian constraints (in particular the one known as Independence of Irrelevant Alternatives) requires the aggregation procedure to respect pairwise judgments in a way that gives rise to our intuitions that the inductive premise is true (cf. Fara's (2000) Similarity Constraint).

An obvious challenge for this kind of analysis is the fact that vagueness is not limited to multidimensional predicates, as we have seen throughout this chapter. (Likewise, vagueness does not disappear if we hold fixed all the dimensions of a multidimensional predicates except for one.) The denotation of *tall* is (arguably, at least) based on a single factor, height, and replacing *healthier* with *taller* in (21) leads to very different judgments: if Kim is taller than Lee, and Lee is taller than Pat, then it absolutely follows that Kim is taller than Pat. Adopting a degree-theoretic approach to gradable adjective meaning of the sort discussed in section 1.3, Grinsell suggests that the locus of multidimensionality in unidimensional adjectives is not the adjective denotation itself, but rather in the "significantly greater than the norm" semantics associated with the positive form of relative adjectives in particular. However, this response relies on a very specific solution to the compositional challenge discussed in section 1.3 — one in which there is a phonologically null morpheme which introduces the relevant multidimensional meaning — which is motivated mainly by theoretical considerations, and not linguistic ones. An intriguing alternative to Grinsell's solution would be to put the "social" back in social choice, and attempt to explain vagueness and the Sorites Paradox not in terms of Arrowian constraints on the denotations of the predicates involved, but rather in terms of Arrowian constraints on the social problem of coordination on uncertain denotations by a community of language users.

References

K. Arrow (1950). 'A Difficulty in the Concept of Social Welfare'. *Journal of Political Economy* **58**(4):328–346.

K. Arrow (1959). 'Rational Choice Functions and Orderings'. *Economica* **26**(102):121–127.

A. Bale (2009). 'A Universal Scale of Comparison'. *Linguistics and Philosophy* **31**(1):1–55.

C. Barker (2002). 'The Dynamics of Vagueness'. *Linguistics and Philosophy* **25**(1):1–36.

R. Bartsch & T. Vennemann (1972). 'The Grammar of Relative Adjectives and Comparison'. *Linguistische Berichte* **20**:19–32.

R. Bartsch & T. Vennemann (1973). *Semantic Structures: A Study in the Relation between Syntax and Semantics*. Athäenum Verlag, Frankfurt.

S. Beck, et al. (2009). 'Crosslinguistic Variation in Comparative Constructions'. *Linguistic Variation Yearbook* **9**:1–66.

M. Bierwisch (1989). 'The Semantics of Gradation'. In M. Bierwisch & E. Lang (eds.), *Dimensional Adjectives*, pp. 71–262. Springer-Verlag, Berlin.

J. Bobaljik (2012). *Universals in Comparative Morphology: Suppletion, superlatives, and the structure of words*. MIT Press, Cambridge, MA.

M. R. Bochnak (2013). *Cross-Linguistic Variation in the Semantics of Comparatives*. Ph.D. thesis, University of Chicago.

A. Bogusławski (1975). 'Measures are Measures: In Defence of the Diversity of Comparatives and Positives'. *Linguistiche Berichte* **36**:1–9.

H. Burnett (2016). *Gradability in Natural Language: Logical and Grammatical Foundations*. Oxford Studies in Semantics and Pragmatics. Oxford University Press, Oxford, UK.

M. J. Cresswell (1976). 'The Semantics of Degree'. In B. Partee (ed.), *Montague Grammar*, pp. 261–292. Academic Press, New York.

D. A. Cruse (1986). *Lexical Semantics*. Cambridge University Press, Cambridge, UK.

D. G. Fara (2000). 'Shifting Sands: An Interest-Relative Theory of Vagueness'. *Philosophical Topics* **20**:45–81.

K. Fine (1975). 'Vagueness, Truth and Logic'. *Synthese* **30**:265–300.

I. Francez & A. Koontz-Garboden (2015). 'Semantic Variation and the Grammar of Property Concepts'. *Language* **91**(3):533–563.

T. Grano & S. Davis (to appear). 'Universal markedness in gradable adjectives revisited: The morpho-semantics of the positive form in Arabic'. *Natural Language and Linguistic Theory* .

T. Grinsell (2017). *Semantic Indecision.* Ph.D. thesis, University of Chicago.

I. Heim (2006). '"Little"'. In M. Gibson & J. Howell (eds.), *Proceedings of Semantics and Linguistic Theory 16*, pp. 17–34. CLC Publications.

H. Kamp (1975). 'Two Theories of Adjectives'. In E. Keenan (ed.), *Formal Semantics of Natural Language*, pp. 123–155. Cambridge University Press, Cambridge. Reprinted in 1984.

C. Kennedy (1999). *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison.* Garland, New York. (1997 UCSC Ph.D Dissertation).

C. Kennedy (2007a). 'Modes of Comparison'. In M. Elliott, J. Kirby, O. Sawada, E. Staraki, & S. Yoon (eds.), *Papers from the 43rd Annual Meeting of the Chicago Linguistic Society Volume 1: The Main Session*, pp. 139–163, Chicago, IL. Chicago Linguistic Society.

C. Kennedy (2007b). 'Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Predicates'. *Linguistics and Philosophy* **30**(1):1–45.

C. Kennedy (2011). 'Vagueness and Comparison'. In P. Egré & N. Klinedinst (eds.), *Vagueness and Language Use*, pp. 73–97. Palgrave, Houndmills, UK.

C. Kennedy & B. Levin (2008). 'Measure of Change: The Adjectival Core of Degree Achievements'. In L. McNally & C. Kennedy (eds.), *Adjectives and Adverbs: Syntax, Semantics and Discourse*, pp. 156–182. Oxford University Press, Oxford, UK.

C. Kennedy & L. McNally (2005). 'Scale Structure and the Semantic Typology of Gradable Predicates'. *Language* **81**(2):345–381.

P. Klecha (2017). 'On Unidirectionality in Imprecision'. *Linguistics and Philosophy* .

E. Klein (1980). 'A Semantics for Positive and Comparative Adjectives'. *Linguistics and Philosophy* **4**:1–45.

E. Klein (1991). 'Comparatives'. In *Semantik: Ein internationales Handuch der zeitgenössischen Forschung* (Semantics: An International Handbook of Contemporary Research), chap. 32, pp. 673–691.

M. Krifka (2007). 'Approximate interpretation of number words: A case for strategic interpretation'. In G. Bouma, I. Krämer, & J. Zwarts (eds.), *Cognitive Foundations of Interpretation*, pp. 111–126. Koninklijke Nederlandse Akademie van Wetenschapen, Amsterdam.

P. Lasersohn (1999). 'Pragmatic Halos'. *Language* **75**(3):522–551.

D. Lassiter & N. Goodman (2014). 'Context, scale structure, and statistics in the interpretation of positive-form adjectives'. In *Proceedings of Semantics and Linguistic Theory*, vol. 24, pp. 587–610.

D. Lassiter & N. D. Goodman (2015). 'Adjectival vagueness in a Bayesian model of interpretation'. *Synthese* .

T. Leffel, et al. (to appear). 'Imprecision Is Pragmatic: Evidence from Referential Processing'. In *Proceedings of Semantics and Linguistic Theory 26*.

D. K. Lewis (1970). 'General Semantics'. *Synthese* **22**:18–67.

S. McConnell-Ginet (1973). *Comparative Constructions in English: A Syntactic and Semantic Analysis*. Ph.D. thesis, University of Rochester.

M. Pinkal (1995). *Logic and Lexicon*. Kluwer, Dordrecht.

C. Potts (2008). 'Interpretive Economy, Schelling Points, and Evolutionary Stability'. Ms., University of Massachusetts, Amherst.

C. Qing & M. Franke (2014a). 'Gradable adjectives, vagueness, and optimal language use: A speaker-oriented model'. In L. Champollion & A. Szabolcsi (eds.), *Proceedings of Semantics and Linguistic Theory 24*, pp. 23–41.

C. Qing & M. Franke (2014b). 'Meaning and Use of Gradable Adjectives: Formal Modeling Meets Empirical Data'. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 36.

C. Rotstein & Y. Winter (2004). 'Total Adjectives vs. Partial Adjectives: Scale Structure and Higher-Order Modifiers'. *Natural Language Semantics* **12**:259–288.

J. Rusiecki (1985). *On Adjectives and Comparison in English*. Longman guistics Library, New York.

E. Sapir (1944). 'Grading: A Study in Semantics'. *Philosophy of Science* **11**:93–116.

G. W. Sassoon (2009). 'Restricted Quantification over Tastes'. In M. Maria AloniHarald BastiaanseTikitu de JagerKatrin Schulz Maria AloniHarald BastiaanseTikitu de JagerKatrin Schulz Aloni, H. Bastiaanse, T. de Jager, & K. Schulz (eds.), *Proceedings of the Amsterdam Colloquium*, pp. 163–172.

G. W. Sassoon (2013). *Vagueness, Gradability and Typicality: The Intepretation of Adjectives and Nouns*, vol. 27 of *Current Research in the Semantics/Pragmatics Interface*. Brill, Leiden.

O. Sawada (2009). 'Pragmatic Aspects of Implicit Comparison: An Economy Based Approach'. *Journal of Pragmatics* **41**:1079–1103.

P. A. Seuren (1973). 'The Comparative'. In F. Kiefer & N. Ruwet (eds.), *Generative Grammar in Europe*, pp. 528–564. Riedel, Dordrecht.

S. Solt (2009). *The Semantics of Adjectives of Quantity*. Ph.D. thesis, The City University of New York.

L. Stassen (1985). *Comparison and Universal Grammar*. Basil Blackwell, Oxford.

K. Syrett (2007). *Learning about the Structure of Scales: Adverbial Modification and the Acquisition of the Semantics of Gradable Adjectives*. Ph.D. thesis, Northwestern University.

K. Syrett, et al. (2010). 'Meaning and Context in Children's Understanding of Gradable Adjectives'. *Journal of Semantics* **27**(1):1–35.

A. Toledo & G. W. Sassoon (2011). 'Absolute vs. Relative Adjectives - Variance Within vs. Between Individuals'. In N. Ashton, A. Chereches, &

D. Lutz (eds.), *Proceedings of Semantics and Linguistic Theory 21*, pp. 135–154.

P. Unger (1975). *Ignorance.* Clarendon Press, Oxford.

J. van Benthem (1982). 'Later than Late: On the Logical Origin of the Temporal Order'. *Pacific Philosophical Quarterly* **63**:193–203.

R. van Rooij (2011a). 'Implicit versus Explicit Comparatives'. In P. Egré & N. Klinedinst (eds.), *Vagueness and Language Use*, pp. 51–72. Palgrave MacMillan, Houndmills, UK.

R. van Rooij (2011b). 'Vagueness and Linguistics'. In G. Ronzitti (ed.), *Vagueness: A Guide*, vol. 19 of *Logic, Epistemology and the Unity of Science*. Springer.

A. von Stechow (1984). 'Comparing Semantic Theories of Comparison'. *Journal of Semantics* **3**:1–77.

A. von Stechow (2009). 'The Temporal Degree Adjectives *früh(er)/spät(er)* 'early(er)'/'late(r)' and the Semantics of the Positive'. In A. Giannakidou & M. Rathert (eds.), *Quantification, Definiteness and Nominalization*, pp. 214–233. Oxford University Press.

S. Wheeler (1972). 'Attributives and Their Modifiers'. *Noûs* **6**(4):310–334.

T. Williamson (1992). 'Vagueness and Ignorance'. *Proceedings of the Aristotelian Society* **66**:145–162.

C. Wright (1975). 'On the Coherence of Vague Predicates'. *Synthese* **30**:325–365.