Vagueness

*Hans Kamp and Galit W. Sassoon**

*Both authors have worked on vagueness on their own for many years at different but extended periods of time. We have both been influenced extensively and decisively by many of our teachers, colleagues and students and much of that – probably much more than we are aware of ourselves - has gone into this chapter, even though its principal task and aim is an even-handed presentation of some of what the rich and diversified literature on vagueness has to offer. We both feel a deep acknowledgement to all whose influence we have consciously or unconsciously profited from, but we see no way of turning this general sense of indebtedness into individual acknowledgements. There are nevertheless a few influences that each of feels clearly and strongly about, that we underwent while the project of writing this chapter was under way and that has had identifiable effects on its outcome and these we do want to mention here explicitly. The first author gratefully acknowledges all he learned from Mark Sainsbury, with whom he jointly taught a graduate seminar in the philosophy department of the University of Texas and subsequently a compact course at NASSLLI in 2012, as well as contributions from the audiences that attended those two events and also from those who were present at a compact seminar at the University of Chicago in the fall of 2012. The second author would like to extend special thanks to Susan Rothstein and to Fred Landman, as well as to acknowledge how grateful she is for all that she had learnt throughout the years from several additional people, including especially, Nirit Kadmon, Frank Veltman and Robert van Rooij.

We also want to thank the editors of this Volume. We originally submitted a draft of this chapter that far exceeded the limits on size that had been communicated to us in advance. In the end it was they who did what should have been our own job, viz. to compress this submission into the present, considerably shorter version. For a somewhat different version of the original submission interested readers may consult the authors' Homepages.

Many thanks for financial support by a Grant from the GIF, the German-Israeli Foundation for Scientific Research and Development.

Vagueness is an ultimate challenge. An enormous diversity of literature on the topic has accumulated over the years, with no hint of a consensus emerging. In this light, Section 1 presents the main aspects of the challenge vagueness poses, focusing on the category of adjectives, and then gives some brief illustrations of the pervasive manifestations of vagueness in grammar.

Section 2 deals with the Sorites paradox, which for many philosophers is the hallmark of vagueness: By assigning a vague predicate step by apparently inescapable step to more and more objects one is eventually led to assign it to entities of which it plainly isn't true.

It is hard to resist the force of the paradox once one has been exposed to it. The result of this has been that many see the philosophical problem presented by vagueness as nothing other than the problem of solving the Sorites. The efforts to solve the Sorites paradox have uncovered a range of important connections between vagueness and other aspects of language and thought. But most of these seem to lead further and further away from what some consider the core issues that vagueness raises.

Given the challenge posed by the Sorites, it is rather remarkable to discover that there is a lot more to vagueness beyond the paradox. In fact, linguists traditionally leave it to the philosophers to deal with the Sorites and put their own efforts into dealing with other manifestations of vagueness in natural language and their consequences for grammar.

Section 3 reviews some of these additional phenomena, centering around three issues: (i) the controversial connections between vagueness and morphological gradability, (ii) the similarity and differences between the phenomena of vagueness and imprecision, and (iii) the ways in which vagueness infiltrates various grammatical constructions we find in language, with consequences for the architecture of grammar. The aim of this section is to highlight the main questions which any theory of vagueness will ultimately have to address.

1. Indications of vagueness

Almost all of the literature on vagueness focuses on the vagueness of predicates. And much of the literature on vague predicates has concentrated on adjectives. For much of this chapter we follow the tradition in both these respects.

1.1 Vagueness in adjectives

Relative adjectives, such as *tall, short, big, small*, *clever* and *obtuse,* have since long been prominent examples of vagueness, as opposed to **sharp adjectives** like *prime* or *even*, which exemplify what it is for an adjective not to be vague. A third type of **absolute adjectives**, such as *clean, dirty, straight* and *bent*, suggests that presence or absence of the features indicative of vagueness may not be a black or white matter (see section 3 for more on this point.) The present section explains this classification.

Positive forms of relative adjectives, as in "Danny is tall" or “The pencil is big”, exhibit a bundle of features indicative of vagueness, as follows.

First, the truth values of predications involving such adjectives are dependent on **the context of utterance**. In particular, they depend on a contextually determined set of entities. Since Klein (1980), this set is called the ***comparison class***. A subset of this class is designated as the contextual extension of the adjectival projection. For instance, a basketball player may count as a tall man, yet not as a tall basketball player; a certain model of a car may count as big for a car, but small for a truck; and an idea can be considered clever for a five year old, but not for a university professor (Kamp and Partee 1995).

Moreover, the usage of relative adjectives involves **a subjective component** (Graff-Fara 2000; Kennedy 2013; Barker 2013; Crespo and Veltman 2014). Decisions about the cutoff between, e.g., tall and not tall entities in a comparison class are affected by the range and distribution of heights of entities in the class. However, there is no uniquely determined method for determining the cutoff even after a comparison class has been fixed (Klein 1980, Kennedy 2007). Decisions may vary depending on speakers’ contextual interests, goals and desires (Graff-Fara 2000). For example, a certain hat may count as *expensive* when it is meant to be worn on a regular basis, but *inexpensive* when you are shopping for something to wear on a special occasion, like your wedding. And even when this goal-related aspect of the context has been fixed, a certain amount of indeterminacy remains: the truth of judgements like *This hat is expensive* or *This boy is clever* can still depend on personal opinion.

In contrast to relative adjectives, absolute adjectives, it has been noted since Unger (1975) and Kennedy (1999), appear to have clear, lexically-determined extensions. For instance, predicating *clean* of a shirt seems to be right only if it doesn’t have a single speck of dirt on it. Otherwise it is dirty, regardless of how it may compare with other entities. Similarly, no comparison class seems needed to determine whether a container is *empty* or a wall is *straight* (Rips and Turnbull

1980), and the truth values of such propositions appear to be immune to interests or desires. Nevertheless, here too we find contextual dependence. The form this dependence takes is indicated by adjective-modifying phrases like *clean for a street in this area* (an example of a **comparison class-determining *for*-phrase;** cf. Cruse 1980) or **judgement-determining phrases** like *in my opinion*, as we find it in *In my opinion, this room is clean*.

Second, relative adjectives clearly exhibit **borderline cases**, for which it is not clear whether they truly apply or not. For instance, out of a group of ordinary pencils that vary in size an 18 cm pencil is obviously big. A 5 cm pencil is obviously not big. But it is not clear whether 10 cm or 12 cm pencils are big or not. Absolute adjectives are different in this respect, at least in first approximation: if cleanliness is a matter of total absence vs. presence of dirt, every entity is either clean or not. But a closer look suggests that here too there is room for uncertainty: When does a shirt truly qualify as clean; when a surface as truly flat?

Let us ignore these more delicate questions for now and assume for the sake of argument that absolute adjectives are sharp, while relative adjectives are not. A simple way to explicate this difference is in terms of the *extensions* of these two types of adjectives. The application domain of an absolute adjective, i.e. the set of things to which the adjective can be applied in principle, is divided into two parts, the (*positive*) *extension* of the adjective, consisting of the entities of which the adjective is definitely true, and the negative extension, consisting of those for which the adjective is definitely false. The positive and the negative extensions cover the entire application domain. In contrast, the application domain of a relative adjective is divided into three parts, a positive extension, a negative extension and a 'truth value gap' consisting of its borderline cases.[1]

Third, some of the context dependence of vague predicates is mediated by **tolerance**: small differences have no consequences for membership (Dummett 1975; Wright 1975; Kamp 1981). In particular, the role of tolerance is easy to recognize in relative adjectives; entities sufficiently resembling extension members count as members too. For example, if a 20 cm pencil is long, then, intuitively, any pencil insignificantly shorter – say, one shorter by just 1 mm – is also long.

It is tolerance that gives rise to **the Sorites paradox**. From intuitively true premises such as (1a-b) we are led in a finite number of steps to the absurd conclusion in (1c).

(1) Mathematical induction Sorites
 a. A 20 cm pencil is long
 b. Any pencil one mm less long than a long pencil is long
 c. A 1 cm pencil is long

This form of the paradox is called the *Mathematical induction sorites*. The paradox can take other forms, all of which must be addressed in a comprehensive theory of vagueness (Sorensen 2013). A second form is (2), the so-called *conditional sorites*, in which the universal premise in (1b) is replaced with a set of conditional premises such as those in (2b,c,d...). Again, the unwarranted conclusion (2e) follows from apparently true premises by seemingly uncontroversial reasoning.

(2) Conditional Sorites

[1] See Kriz and Chemla for a discussion of experimental methods for research of truth value gaps.

a. A 20 cm pencil is long
b. If 20 cm pencil is long, then a 19.9 cm pencil is long
c. If 19.9 cm pencil is long, then a 19.8 cm pencil is long
d. …
e. 1 cm pencil is long

Note that the only inference principle that is needed to derive the conclusion of (2) from its premises is Modus Ponens (MP). MP also plays a central part in the most straightforward derivation of the conclusion of (1). There are other formulations of Sorites arguments, which suggest the application of other inference principles than MP. But we do not believe that from a proof-theoretic point of view those alternative forms raise issues that are significantly different from the ones that arise in connection with the role of MP in deriving the conclusions of (1) and (2). In our discussions of Sorites arguments in Section 2, we will, therefore, focus on arguments of these two forms and on the role of MP in obtaining their conclusions.

There is one other variant of the Sorites paradox, however, not involving MP, that we do want to mention explicitly, because it brings out the conceptual problems connected with the paradox in a way that (1) and (2) do not. (3) is like (1), but with the second premise and conclusion contraposed. The conclusion is now that there exists a cut off point between the positive and the negative extension of the predicate 'long pencil'. That, the argument is trying to tell us, follows from the fact that there exist pencils that are definite instances of the predicate and also pencils that definitely do not instantiate the predicate.

(3) Line drawing Sorites
a. A 200 millimeters pencil is long
b. For some number n of millimeters, it is not the case that an n millimeters long pencil is long.
c. There is some number n equal or smaller than 200, such that an n millimeter long pencil is long, while an n – 1 millimeter long pencil is not long.

On a first pass, absolute adjectives are immune to the paradox. However, we already noted that absolute adjectives behave as sharp predicates only up to a point and that the picture is, therefore, more complex. This also applies to many non-adjectival predicates that may appear sharp at first sight, but where a closer look reveals fuzziness at the fringes. To see this, let us examine two important classes of absolute adjectives and compare them more closely to relative adjectives. Absolute and relative adjectives have been claimed to differ in the type of **modifiers** they license (see discussion in McNally, this volume). Modifiers such as *slightly* and *perfectly* appear to favour predicates for which a conventional criterion exists to determine membership. More specifically, they have been analyzed as relating to membership criteria defined in terms of minimum and maximum membership thresholds. For instance, the modifier *slightly* felicitously combines with adjectives like *dirty* whose membership criterion merely requires the existence of some minimal amount of dirt, and *perfectly* combines with adjectives like *clean* whose criterion requires maximum cleanliness or complete absence of dirt. This distinction between adjectives governed by minimum criteria and adjectives governed by maximum criteria seems to explain the contrasts between (4a,b) and (4c) (Kennedy & McNally 2005; Kennedy 2007).

(4) Degree Modifiers indicating presence/absence of vagueness

a. {slightly, #perfectly} {dirty, wet}
b. {#slightly, perfectly} {clean, dry}
c. {#slightly, #perfectly} {tall, short}

The acceptability of this type of modifiers with relative adjectives is reduced, as in (1c), and becomes fully acceptable only when an appropriate membership criterion is contextually accommodated, for instance, when, e.g., *slightly tall* is interpreted as conveying 'slightly too tall for some (contextually determined) purpose', or, alternatively, when the expression is used to refer to borderline cases, which can be thought of as constituting the fuzzy boundary between positive and negative extensions (Kagan and Alexeyenko 2011; Sassoon 2012; Solt 2012; Lassiter 2011). Relative adjectives co-occur more frequently and naturally with 'strengthening' modifiers like *very* or *extremely* (Syrett 2007).

Another distinction between absolute and relative adjectives comes from **inference patterns** such as those illustrated in (5a-c) (Rotstein and Winter 2004; Kennedy & McNally 2005). The distinction relates to different membership criteria. The comparative "X is dirtier than Y" entails that X is covered by some amount of dirt and, therefore, counts as dirty. Similarly, the comparative "X is cleaner than Y" entails that Y necessarily fails to be maximally clean, and therefore counts as not-clean. By contrast, a relative adjective selects a contextual standard; therefore a comparative form such as "X is taller than Y" does not entail that either entity is in the positive or in the negative extension.

(5) Inference patterns indicating presence/absence of vagueness:

a. Minimum criterion:	X is dirtier than Y	$\Rightarrow$	X is dirty.
b. Maximum-criterion:	X is cleaner than Y	$\Rightarrow$	Y is not (maximally) clean.
c. Contextual criterion:	X is taller than Y	$\not\Rightarrow$	X is tall / Y is not tall.

Back to tolerance, Burnett (2014a,b) argues that features of vagueness can be manifested **asymmetrically** in absolute adjectives with a minimum- and maximum-criterion. Recall that when one mm counts as an irrelevant difference, then both *tall* and *not tall* are tolerant, namely speakers assent to both (6a) and (6b), respectively, where $\rightarrow$ stands for material implication.[2] By contrast, both the sharply defined predicate *prime* and its negated form *not prime* are intolerant, which is shown by the fact that speakers assent to neither (6c) nor (6d).

(6) Symmetric Vagueness: For all x; y:

a. x is tall and x and y's heights differ by a millimeter $\rightarrow$ y is tall.	(true)
b. x is not tall and x and y's heights differ by a millimeter $\rightarrow$ y is not tall.	(true)
c. x is prime and x and y differ by one $\rightarrow$ y is prime.	(false)
d. x is not prime and x and y differ by one $\rightarrow$ y is not prime.	(false)

For maximum-criterion absolute adjectives, such as *straight*, the matter is different. Here the positive form may be tolerant – that is, in at least some contexts speakers assent to (7a) – but at the same time, the negated form is not tolerant; e.g., (7b) is falsified by individuals x with a single 1mm bend, which may count as not straight, and individuals y that have absolutely no bends and which, therefore, cannot but count as straight. In minimum-criterion absolute adjectives such as *bent*, the negated form may be tolerant – in at least some contexts, speakers

[2] A natural way to express (6a) is: "Every individual y one millimeter shorter than a tall individual x is tall as well".

assent to (7d) – but the positive form cannot; for example, (7c) is falsified by a rod x with a single 1mm bend, which can be considered bent, and a rod y that has absolutely no bends and, therefore, does not qualify as bent.

(7) Asymmetric Vagueness: For all x; y:
a. x is straight and x and y differ by a 1mm bend → y is straight. (sometimes true)
b. x is not straight and x and y differ by a 1mm bend → y is not straight. (always false)
c. x is bent, and x and y differ by 1mm bend → y is bent. (always false)
d. x is not bent, and x and y differ by 1mm bend → y is not bent. (sometimes true)

The contexts in which speakers assent to (7a) or (7d) are contexts in which they are willing to consider as *straight* or *not bent* also objects which are not maximally straight. For these expressions, **there are some contexts** in which they give rise to borderline cases, fuzzy boundaries and Sorites arguments. Burnett, who demonstrates this in detail, calls these expressions 'potentially vague'. There seems to be a great deal of variation between speakers on this point. Some accept tolerance-governed interpretations of absolute adjectives, others do not. This variation is reflected in a debate among scholars over whether these interpretations are manifestations of vagueness or of mere imprecision (cf. section 3.2).

Fourth, it is important to distinguish between the theoretical problem that tolerance-governed predicates present because of their occurrence in Sorites arguments and the practical constraints that tolerance may be thought to impose on the actual use of such predicates: any given context in which the predicate can be used should restrict the set of entities to which the predicate is applied in such a way that somewhere there is a 'jump' within the set: a partition of the set into two non-empty subsets S^+ and S^-, such that no member of S^- falls within **the tolerance range** of any of the elements of S^+. If such a jump separates the clear positive cases and clear negative cases of the tolerance-governed predicate, no Sorites argument can be constructed for it on the basis of the given set; and the Sorites problem can be kept at bay at the conceptual level by assuming that there exists a partition separating positive and negative extensions and that the boundary between positive and negative extension runs somewhere through the middle of the gap (Pagin 2010a,b; Gaifman 2010; van Rooij 2011; Cobreros et al. 2011; Burnett 2011).

One aspect of the context dependence of tolerance-governed predicates is the tolerance relation itself. How similar must two entities be to count as within each other's tolerance range? This often depends on context as well. For instance, as noted by Graff-Fara (2000), it depends on the context whether 'one grain of coffee makes a difference' – it doesn't when you make a pot of coffee in the way we normally do that, but it may in the context of a scientific experiment (e.g. one which investigates how different amounts of caffeine affect mice). Translated into the language of tolerance-governed predicates: for any number k, k and k+1 grains may be within each other's tolerance range for the predicate 'a suitable number of grains of coffee' in one context of use, but not in some other context.

Fifth, there is the question of how vagueness relates to **ambiguity**. (8) shows that the ellipsis test for ambiguity (Zwicky and Sadock 1975; Barker 2006) works out differently for vague predicates than it does for classical examples of ambiguity like the noun *bank* or the verb *wave* (ambiguous between the waving done be people and the waving done by flags).

(8) a. Ann went to a bank, and Bill did too.
b. #Ann waved, and the flag did too.
c. Bill is tall, and his six year old daughter is too.

(8a) has only two readings: (i) Ann went to a financial institution and Bill went to a financial institution, and (ii) Ann went to a river bank and Bill went to a river bank. This indicates that the ambiguity of bank in the first conjunct must be resolved before the first conjunct can be used to interpret the second, elliptical conjunct. (8b) is another illustration of the same point. Here the subject of the first clause, a person, disambiguates the occurrence of *wave* in this clause to the sort of waving that is done by people. Once that disambiguation has taken place, the first clause is no longer suited to resolve the ellipsis of the second clause. If the adjective *tall* was ambiguous in this same sense between 'tall as applied to human adults' and 'tall as applied to human six year olds' then (8c) should have been false (or else the description of an anomalous situation in which Bill's daughter is some giant monster). But (8c) can be truly asserted of a perfectly normal situation, in which Bill is tall for an adult and his daughter tall for a six year old girl.

A similar difference arises between (9a) and (9b) with regard to the interpretations of *bank* and *tall* as applied to different members of the quantifier restriction (Kennedy 2007).

(9) a. Everyone in Bill's family went to the bank.
b. Everyone in Bill's family is tall.

Sixth, an important aspect of relative adjectives is their behavior in compound predications that have the logical form of classical contradictions or tautologies. It was long held that, for instance, the predicate 'tall or not tall' was tautological, and could therefore be regarded as true of all entities in *tall*'s application domain, including its truth value gap; and that, likewise, the predicate 'tall and not tall' is false of all entities in the application domain (Kamp 1975; Kamp and Partee 1995; Barker 2006). That compound predicates like 'tall or not tall' and 'tall and not tall' can be seen as true or false of entities in the truth value gap of *tall* has been seen as evidence for the **Supervaluation** approach to vagueness (see Section 2.2.2 below). But more recent experimental work (Ripley 2011; Serchuk et al. 2011; Alxatib and Pelletier 2011; Egre, Gardelle and Ripley 2013) casts doubts on this view. Rather than making vacuous claims, predications like '*a* is tall or not tall' and '*a* is tall and not tall' are used to state or highlight that *a* belongs to the truth value gap of *tall* and **not** to its positive or negative extension.

One problem with these results is that there is considerable variation among speakers' reactions to clauses of the form '*a* is tall and not tall'. There have been different responses to these facts in the literature. We mention two of these:

(i) When we interpret the utterance of other speakers, we normally don't assume that they knowingly make empty or contradictory claims. So when it superficially looks like that is what they are doing, then we try to reinterpret what they say. This is how we deal in particular with utterances of apparent contradictions and tautologies like '*a* is tall and not tall' and '*a* is tall or not tall': We try to reinterpret the different occurrences of 'tall' they contain in such a way that their content becomes non-vacuous (i.e. non-empty and non-contradictory; Kamp and Partee 1995).

(ii) The data involving apparent tautologies and contradictions with vague predicates show that one or more of the logical words in these predicates – *and, or, not* – are used in a sense that differs from the ones attributed to them in the semantics of classical logic. The task for the theorist, therefore, is to identify what the different sense or senses are and to determine how these affect the logic of *not*, *and* and *or* (Ripley 2011; Serchuk et al. 2011; Alxatib and Pelletier 2011; van Rooij, Cobreros, Egré and Ripley 2012; Burnett 2014b).

At the present time, no consensus has been reached as to which of these two strategies has got things right. Perhaps different strategies are right in capturing the linguistic and logical dispositions of different groups of speakers.

1.2 Beyond adjectives

While adjectives seem especially prone to truth value gaps and to the contextual variability that is indicative of vagueness, it is clear that vagueness is not limited to this category. Widely discussed examples of vague nouns are *heap* and *mountain*, which have played a key role in the literature on predicates that give rise to borderline cases, fuzzy boundaries and Sorites problems. Intuitively, a collection of one or two grains does not make a heap, and if no collection of n grains makes a heap, neither does any collection of n+1 grains. The paradoxical conclusion is that for no n does a collection of n grains make a heap. Much the same goes for *mountain*. If an elevation of 2500 meters counts as a mountain, how could an elevation of 2499 meters fail to qualify? And where is the dividing line between mountains and hills?

Tolerant scalar properties such as number of grains, or merely size or volume, appear to play a role in classification under these nouns. They help determine what constitutes a heap or a mountain, although other properties, such as substance and shape, matter as well. Similarly, classification of an individual *a* under nouns like *child*, *adult*, *boy* or *girl* depends in part on how old *a* is, and the vagueness of *old* infiltrates the semantics of these names as well. (A noun is vague as soon as at least one of its satisfaction criteria is vague.)

Typically, nouns – for example, *chair* or *tomato*, but the choice is quite arbitrary – are multidimensional (Kamp 1975) and relatively sharp (Kamp and Partee 1995). Psychological research, however, suggests that nouns do have a gap similar to that seen in vague adjectives. For example, tomatoes fall somewhere between fruit and vegetables and three-legged seats with a small back fall between chairs and stools (Hampton 1997; Labov 1973). Speakers rarely change their minds about the categorization of clear instances of nouns, but they do so often in borderline cases like curtains for *furniture* or avocado for *vegetables* (McCloskey and Glucksberg 1978).

Moreover, classification under nouns is governed by similarity. Entities classify under nouns if their scores along multiple dimensions sufficiently match the ideal values for the noun, or if they are closer to the ideal values of the noun than to the ideal values of any salient alternative (Gärdenfors 2004; Hampton 1998). van Deemter (2010) provides ample illustrations of the way similarity and tolerance join to give rise to the Sorites paradox in the sharpest of nouns, including natural kind nouns like *frog* or *tadpole*, or mathematic concepts like *theorem, corollary*, or *finitism.*

Importantly, expert taxonomical criteria are vulnerable to vagueness. For example, the traditional foundation for the division of the animal kingdom is interbreeding: two animal groups belong to the same species if they can interbreed. But in actual fact the interbreeding criterion is far less unequivocal than must have been thought when it was first introduced, and when taken strictly, it leads to a gamut of mutually overlapping species (see van Deemter 2010, chapter 2, for a discussion of the insufficiency of this and other scientific criteria for natural kind terms.)

These various manifestations of vagueness strongly suggest that nouns are vaguer than is usually acknowledged in the linguistic literature. Like vague adjectives, nouns provide evidence for borderline cases, contextual variance in categorization judgments, lack of precise boundaries and the Sorites paradox.

In fact, this aspect of vagueness penetrates into the sharpest looking words. The *meter*, for example, is defined as the distance between two marks on a certain rod in Paris, when found at a certain temperature and atmospheric pressure. However, as Russell (1923) points out, the marks are not points, so the distance between them is not defined with total precision. In addition, temperature is never quite uniform within an object. And so it is with physical magnitudes other than length: There is always some margin of error involved in their measurement. So vagueness encroaches even upon physical predicates that are given by what look like perfectly sharp definitions (cf. Williamson 1994). Moreover, these considerations pertain, mutatis mutandis, to verbs as well as to nouns (see, for instance, Altshuler and Schwarzschild 2014 for a discussion of fuzzy boundaries in stative predicates in relation to the cessation problem).

It has been argued that even proper names are subject to vagueness (Russell 1923; Lewis 1988). Typically, the question when a proper name refers to its bearer has no fully precise answer. There may be long stretches of time when a name attaches firmly and unequivocally to its bearer; but there are, nevertheless, times when this is not clear, e.g. at the time of birth of a person or of the creation of an artifact or the foundation of a city; and likewise at the end of an entity's lifespan – the death of the person, or the destruction of the city or artefact. In the case of proper names, however, it is less clear that we must situate the vagueness in the name, rather than seeing it is a kind of vagueness or indeterminacy in the temporal extent of the thing named. This second possibility of vagueness in the temporal or spatial extension of real entities, also known as ' vagueness in the world' or 'metaphysical' or 'ontological' vagueness, has long been a topic of dispute within philosophy. Some philosophers acknowledge vagueness in the world as *sui generis,* while others follow the early lead of Russell in regarding vagueness exclusively as a matter of language (or at most as a matter of general cognition). Since the focus of this article, as part of a handbook on Semantics, is on vagueness in language, vagueness in the world is not one of its themes and we will, therefore, leave it at this brief mention of it. But some of the issues raised by the possibility of metaphysical vagueness are of the first importance, and in a more broadly conceived discussion of vagueness than would be appropriate here the topic would deserve careful discussion (Unger 1980; Tye 2001; Van Inwagen 1998).[3]

Other domains in which vagueness is prominent are those of quantifying expressions, frequency adverbs (Lewis 1979) and adjectival modifiers (Lakoff 1987). For example, *many* and *few,* as they are found, for instance, in statements of the form "Many/few As are Bs", are evidently vague, while expressions such as *every, some,* and *three* have generally been treated as sharp. Similarly, *very* and *slightly* are vague modifiers, while *perfectly* and *completely* may be regarded as sharp (Lasersohn 1999). *Usually*, *generally*, and *normally* are vague quantificational adverbs, while *always, sometimes,* and *three times* are sharp (Kadmon and Landman 1993; Cohen, this volume; Veltman 1996). The quantifier *Most* is vague, while *more than half* and *more than two thirds* are sharp (see Solt 2014b for significant corpus results and discussion). A notorious example of vague grammatical constructions are those expressing genericity, among them generic indefinites (Carlson and Pelletier 1995)

[3] The focus of this chapter is on semantics and the mytaphysical aspects of vagueness are not directly relevant for the semantic issues we want to discuss. Mytaphysical vagueness will therefore not be considered here.

Notice, however, that sharp quantifiers or modifiers like *every*, *ten* or *one hundred* are implicated in two other kinds of vagueness. First, like all quantifying expressions, they usually involve *domain restrictions* (von Fintel 1994). Sometimes, these are fixed explicitly and precisely, as in 'Every egg/ten eggs in this basket is/are spoiled'. But often domain restrictions are determined by context, and often only loosely, thereby making room for vagueness in the semantic contribution of the quantifying expression. Second, an expression like *everyone* is often used in a way that is loose in another sense. 'Everyone is coming to the meeting' may be accepted as true even if a member of the council who has been in a coma for some time will (obviously) not be coming. And 'One hundred people turned up at the lecture' will often be accepted as correct, although there were in fact only 99 or 98 (Lasersohn 1999). Note that this kind of sloppy use of round numbers seems to implicate vagueness of the kind governed by tolerance: if the presence of 98 people at the meeting is enough to justify the claim that a hundred turned up, what about 97? And if 97 is enough, what about 96? Likewise, if the absence of the council member who is in a coma doesn't falsify 'Everyone is coming to the meeting', what about some other council member who is in good health, but who has never attended any of the meetings in the past and has made a point of declaring openly that he will never go? (For more on this see Section 3.2)

1.3. Intermediate summary and a look ahead

In this first perusal of our topic, with its primary focus on relative adjectives as prototypical examples of vague predicates, we have come across the following features that are indicative of vagueness and that affect the truth conditions of vague predications:

(10) Indicators of vague expressions
  a. Contextual variability in truth value judgments
  b. Dependence on a comparison class
  c. A subjective component in truth value assessment.
  d. Borderline cases
  e. Fuzzy boundaries
  f. Tolerance
  g. Susceptibility to the Sorites paradox
  h. Behavior under ellipsis
  i. Informative uses of contradictions and tautologies
  j. Modifiers that select for certain types of vague predicates

All these features raise doubt about the logic of vague predicates. In particular, a potential threat is posed by (10d). Classical logic is the logic of sharp predicates: it is the contribution made by sharp predicates to the truth conditions of the sentences they are part of that ultimately determines the relation of logical consequence (i.e. the semantic characterisation of logical validity) to be that of the classical predicate calculus. It may not be immediately obvious what happens when partial predicates are taken into account as well as sharp ones, but it seems reasonable to suspect that this will affect the consequence relation to the point of changing the logic.

The question what effects vagueness has on logic has led to a multitude of different reactions and to proposals of a considerable range of different vagueness logics. A proper survey of these proposals would require an article of its own. In a handbook on semantics, the logical aspects of vagueness are a lesser priority. Nevertheless, if we see it as part of the task of semantics to explain how the linguistic form of a sentence determines the information from which an interpreter can proceed to derive further conclusions, through the application of certain logical principles, then the question what these principles are, and what the logical form of the information is to which they are applied, cannot be ignored. In what follows in the next sections, we have tried to strike a balance between aspects of logic that are of intrinsic importance to someone with a central interest in the logic of vagueness and the relevance that logical questions have for the semantics of vagueness. The logic of vagueness will be the central concern of section 2. We conclude section 1 with a few preliminary observations.

An early view of the relation between vagueness and logic was thoroughly pessimistic. Frege (1903), Russell (1923) and Dummett (1975) have argued that there can be no coherent logic of vague expressions. The methodological consequences of such a view are daunting. On the plausible assumption that truly sharp predicates can be found only within pure mathematics, this view entails that pure mathematics is the only domain where formal logic (as it was established, first and foremost, by Frege and Russell) has any application. Everywhere else the use of language is essentially incoherent and any attempt to apply formal logic pointless (or worse).

Not only is such a position deeply unsatisfying, it seems to be demonstrably false. Consider for instance the following argument (Keefe 2007).

(11) Unproblematic inference patterns with vague expressions

a. Premise 1: Fred has less than 500 hairs.
b. Premise 2: Anyone with less than 500 hairs is bald.
c. Conclusion: Fred is bald.

There is a strong intuition that this argument is valid, notwithstanding the occurrence in it of the evidently vague predicate *bald*, and that it is valid for exactly the same reasons that would apply to an argument of the same form in which *bald* is replaced by some sharp predicate. Surely, it ought to be possible to say something about the logic of vague predicates that does justice to this intuition, and to formulate and justify a vagueness logic in which (11) is valid.

There is another position, diametrically opposed to that of Frege, Russell and Dummett, according to which this is straightforward. According to this position vagueness is just a form of ignorance (Williamson 1994). When an entity *a* appears to us to belong to the truth value gap of the 'vague' predicate P, that is always because we do not *know* whether or not it is a case of P. Some of that knowledge is irremediable, because of the way in which we are situated, as members of a speech community who have to acquire their language as well as they can, and who cannot do better than use predicates like P to the best of their knowledge of how the satisfaction conditions of P are grounded in the language that is the common property of their community. In that sense, the borderline cases of P may have a kind of absolute status for the language user. But according to the *epistemicist* view of vagueness (the one we are discussing) this epistemic fact about the extension of P must be distinguished from the true (though unknowable) satisfaction conditions of P. These are the satisfaction conditions of a sharp predicate and it is they that enter into the definition of logical consequence, which will therefore give us the classical notion.

The ignorance of the user of P, to whom it looks as if P has borderline cases and that *a* is one of them, can be analysed further within the setting of **epistemic logic**, in which, for any agent *u*,

K_u is a modal operator representing *u*'s knowledge (so $K_uP(a)$ can be read as '*u* knows that *a* has P' and $\neg K_uP(a)$ & $\neg K_u\neg P(a)$ as '*u* doesn't know whether or not *a* has P'.) The semantics of K_u can be given in the form of a set of complete models (i.e. models in which all predicates have gap-free extensions). Among these models, there will be some in which P has its 'true' extension; but there will also be other models, in which P has different extensions. Thus, the totality of the set of models reflects, via the spread of extensions for P represented within it, *u*'s lack of knowledge about P – *u* does not know which model in the set is the one that represents the way things actually are; her knowledge about P is restricted to those parts of its positive and negative extension that are shared between all models in the set.

A third approach – the so-called **supervaluation** approach – takes ostensibly vague predicates P to be genuinely vague, and models genuine vagueness of P as P having a partial semantics, consisting of positive extension, negative extension and (non-empty) truth value gap. But the 'true' partial model M_0, in which each such predicate P may have a tripartite extension, is seen in this approach as part of a more complex structure in which there is also a set of complete models extending M_0, in each of which the truth value gaps of vague predicates in M_0 are resolved.

In these ways, both the epistemicist and the supervaluational approach lead to sets of complete models. But the conceptual motivations behind these sets are very different. And because of that the role that these sets play in what the two approaches have to say about logic are quite different too. As we just observed, for the epistemicist, these sets are relevant only for the analysis of the users' knowledge about P. In contrast, for the supervaluationalist, the set of complete models is relevant to what is true about P, and as we will see in the next section, it plays a decisive part in characterizing the basic logic of vague predicates (and not just the epistemic logic that is obtained by adding the operators K_u). We will see that one way in which the supervaluationist can exploit the set of complete models recovers classical logic (thus reaching the same conclusion as the epistemicist, albeit by a different route). But for the supervaluation approach there are other ways of characterising the logic as well and those lead to non-classical logics.

In addition to these three perspectives on the semantics and logic of vagueness there are also a number of others, which acknowledge that vague predicates have borderline cases, but which do not see a borderline case *a* of P as one where the predication P(*a*) lacks a truth value, but in which the truth value of P(*a*) is an *intermediate* one, somewhere in between the extreme values of 1 (for definite truth) and 0 (for definite falsehood). There are different ways in which this idea can be implemented, from proposals that work with just one intermediate value (and thus with three truth values *in toto*) to those which assume that the possible intermediate values cover the entire open interval (0,1) of real numbers between 0 and 1 (so the entire truth value space is the closed interval [0,1], with, once again, the extreme values 0 and 1 representing definite falsehood and definite truth; Zadeh 1965; Goguen 1968; Edgington 1997). Adopting intermediate truth values leads more or less inevitably to the adoption of a different semantics for the basic logical operators (specifically, *not*, *and*, *or*, *if,.. then*, and the existential and universal quantifier). And the changes in the semantics of these notions lead to different, **non-classical logics**.

A multi-valued approach has also been favoured for an entirely different reason. Here the basic motivation is that, when *a* is a borderline case of P, this doesn't necessarily mean that P(*a*) is neither true nor false, but that P(*a*) is **both** true **and** false: vague predicates have truth value 'gluts' rather than truth value gaps. One consequence of this view is that when *a* is a borderline case of P, P(*a*) and ¬P(*a*) can be both true. This result has been seen as corroborated by the fact mentioned earlier that speakers sometimes use sentences of the form '*a* is P and it is not P' as a way of indicating that *a* is a borderline case of P. This last approach to vagueness is a special case of a more general approach to logic known as **paraconsistency**. Paraconsistent logics are weaker than classical logic and block in particular the classical inference principle *Ex falso sequitur quodlibet*, which allows the inference of any conclusion from a contradictory set of premises (Hyde 1997; Priest 2002; Weber 2010; Ripley 2011).

One final complication of the already complex landscape of vagueness logics has come into focus with the work of Cobreros, Egré, Ripley and Van Rooij (Cobreros et al. 2010; Cobreros et al. 2012). This approach makes use of classical models, in which vague predicates have complete extensions, but in which there is also a (complete) extension for a tolerance relation ~P for a vague predicate P (where *a* ~P *b* means that '*a* and *b* are not significantly different with regard to P'.) These classical models are used to define truth for sentences containing P in three different ways (tolerant, classical, and strict). This creates a number of different possibilities for defining logical consequence as preservation of truth, viz. by requiring possibly distinct notions of truth for the premises and conclusion of a consequence relation.

1.4. Something that will not be discussed in this chapter: Higher Order Vagueness

It has been argued that modelling the vagueness of a predicate as tripartite division of its application domain into positive extension, negative extension and truth value gap is too simple. Intuitively, the boundaries between the three parts of the application domain – e.g., the boundary between the definitely big pencils and the borderline big pencils, and that between the borderline big and definitely not big pencils – cannot be sharp. They too should be fuzzy in their turn, like regions populated by their own, 'second order' borderline cases. But how should we model this? Suppose that there is a (second order) borderline region between the positive extension and the truth value gap of 'big pencil', and another region between the truth value gap and the negative extension. What are we to say about the boundaries between the five regions that this leads to? Consider for instance the boundary between the positive extension and the borderline area between it and the original truth value gap. Do we have any better reason to think that this boundary is sharp? And if not, and we try to do justice to the notion that this boundary is really also a smear of borderline cases, of yet a higher degree, are we to acknowledge even more regions? There is evidently no natural end to this. Each time we refuse to accept that a boundary is sharp and assume that it is really a region of sorts, we will encounter the same problem at the next level.

Fuzziness of first order boundaries is known as **second order vagueness**. Likewise, fuzziness of the boundaries at the next level is referred to as **third order vagueness**, and so on. Vagueness anywhere from second order up is subsumed under the term '**higher order vagueness'**.

Higher order vagueness is perhaps the single most controversial aspect discussed within the vagueness literature. The variation in opinions on this topic is extreme. Here are some of the views that have been expressed:

(i) the view that higher order vagueness exists and that its manifestations are as real, and as essential to the nature of vagueness as its first order manifestations and that no analysis can therefore be complete unless it deals with higher order vagueness too (Sorensen 2010)

(ii) the view that higher order vagueness is an illusion and that those who endorse the notion are driven to it only because their approach to the problems of vagueness gets off on the wrong foot to begin with (e.g. Wright 2012).

(iii) the view that what may look at first blush like second order vagueness in P on closer inspection always turns out to be first order vagueness in some other predicate or predicates (Shapiro 2005).

In the vagueness literature higher order vagueness is often discussed in conjunction with operators, represented as Def (for 'definitely') or Δ, which make it possible to express semantic properties like partiality of vague predicates in a language that also contains those predicates themselves. The first purpose for which such operators are wanted is often to express the absence of truth value. For instance, when *a* is in the truth value gap of P, then a special device is needed to express this in the language of P. ¬P(*a*) obviously won't do, for (on the most natural interpretation of ¬ and one that we do not want to do without) that is the claim that *a* belongs to the negative extension of P, and not to the truth value gap. Given the intended interpretation of Def, an intuitively plausible way of saying that *a* is a clear case of P is 'Def P(*a*)' and, likewise, 'Def¬P(*a*)' would be an intuitively plausible way of saying that *a* is clearly a case of not-P. So, by the same token, a plausible way to say that *a* is in the truth value gap of P is the formula '¬Def P(*a*) & ¬Def¬P(*a*)'.

But an operator serving these purposes has to be handled with care. One question that needs careful consideration is the relationship between Def P(*a*) and P(*a*). Both formulas seem to express the fact that *a* belongs to the positive extension of P. But they cannot be equivalent in the sense that they can be substituted for each other without change of truth conditions. For such an equivalence between Def P(*a*) and P(*a*) would make ¬Def P(*a*) & ¬Def¬P(*a*) equivalent to ¬P(*a*) & ¬¬P(*a*); and unless other changes are made, ¬P(*a*) & ¬¬P(*a*) is a contradiction, and so it cannot express that *a* is in the truth value gap of P.

Def is also often used to express distinctions that have to do with higher order vagueness, for instance, the formula Def φ & ¬Def Def φ is used to indicate that φ is first order but not second order definite. Exactly what such iterations of Def come to depends on logical and semantic properties of the language as a whole to which Def belongs as well as on the particular role that Def plays within it.

These observations indicate that the logic and semantics of Def and of any language in which it occurs side by side with vague predicates are a delicate matter. If the main focus of this chapter had been on the logical and philosophical aspects of vagueness, then a proper discussion of higher order vagueness and operators like Def would be a must. But since our primary focus is on the semantics of vagueness as it manifests itself in the ways we speak, and since there isn't room to discuss all aspects of vagueness, we have decided not to include higher order vagueness among the topics of this contribution.[4]

2. The Sorites paradox

In Section 1 we presented three versions of the Sorites Paradox. In this section we start with responses to version (1) of the argument, which we repeat in (12), this time with the number predicate 'is a number n such that n grains of sand can make a heap' replacing the predicate 'long pencil'. Later on in this section, we will have quite a bit to say about the argument form in (2), of which (13) is a restatement, using the same number predicate.

[4] Some important contributions to the discussion over higher order vagueness not mentioned in the text include Fine (1975), Williamson (1994, 1999) and Barker (2006).

(12) a. Base premise: 1,000,000 grains of sand can make a heap of sand. (Premise 1)
b. Inductive premise: (For all n) If n+1 grains of sand can make a heap then so can n grains. (Premise 2)
c. Conclusion: One grain of sand can make a heap.

(13) a. Base premise: 1,000,000 grains of sand can make a heap.
Conditinal premises
b. If 1,000,000 grains of sand can make a heap, then so can 999,999 grains.
c. If 999,999 grains of sand can make a heap, then so can 999,998 grains.
.
.
d. Conclusion: I grain of sand can make a heap.

The responses are discussed in the following sections, divided into the following types (Sorensen 2013):

i) Denial of the first premise of one version of the paradox (skepticism).
ii) Denial of the second premise (epistemicism, supervaluationism, subvaluationism).
iii) Constraints on the usage of vague predicates (pragmatic contextualism)
iv) Denial of the validity of the argument (multiple truth values/notions).

2.1 Denial of the first premise: Skepticism

The first type of response denies its base premise. This is one way to remove the paradoxality from the argument: If one of its premises is false there is nothing paradoxical about deriving a false conclusion from it. But is there anything else that speaks in favour of this approach? Of course we can infer that the first premise is false on the assumption that the second premise is true, the conclusion is false and the argument valid. But as an explanation of the paradox that obviously won't do. What we need are reasons why the first premise is false that are independent of those assumptions.

But what reasons could there be? You are looking at a heap of 1,000,000 grains of sand. It is clearly and unambiguously a heap, by any standards for the application of *heap*. How is anybody going to argue you out of that, if it isn't by arguing that the concept is paradoxical, or incoherent in a way that the Sorites paradox reveals? If there are really no other ways of resolving the paradox – no way of faulting the other premise or the logical validity of (12) – then, perhaps, you might give up your resistance to this conclusion. But other avenues should be explored before we succumb to something quite as drastic.

Furthermore, let us, for the sake of argument, go along with the assumption that the argument is valid and its second premise true. Why, then, should we infer that the premise must be false rather than that the conclusion is true? Either inference flies in the face of what appear to be incontrovertible facts. One seems to be as bad as the other. Why choose between them?

We can avoid choosing as follows. What follows from the two assumptions we have retained is that the first premise and the conclusion of (12) have the same truth value.[5] If we assume in

[5] This is not quite correct as it stands. An additional principle is needed to exclude the possibility that the first premise is false and the conclusion true. For sorites arguments involving number predicates this principle is a monotonicity principle to the effect that when n < m and n satisfies the predicate then so does m. For other predicates, such as e.g. *big pencil*, the monotonicity principle has to be stated somewhat differently, but the general idea should be clear.

addition that for any two elements *a* and *b* from the application range of the predicate the construction of a sorites sequence is a conceptual possibility, thereby linking *a* and *b* in the manner that 1,000,000 and 1 are connected in argument (13), then it follows that the positive extension of the predicate is either empty or all of its application domain. In other words, either no number n is such that n grains of sand can make a heap or else all numbers n have this property.

This is still strong medicine. But there are some predicates for which the 'all-or-nothing' conclusion involved – either nothing satisfies the predicate or else everything does – does have some independent support. We saw in Section 1 that it is necessary to distinguish between relative adjectives and absolute adjectives and that the latter come in two varieties, maximum criterion (*clean*, *flat*) and minimum criterion (*dirty*, *uneven*). The strict satisfaction conditions for a maximum criterion adjective like *flat* are to the effect that a surface x is flat if and only if x has no unevennesses: as soon as there is some bit of unevenness to x we can imagine a surface y which is just like x except flatter. But that would disqualify x as flat.

If this is the right definition of *flat*, then there are arguably no flat surfaces in physical reality. And if that is true, then *flat* has an empty extension in the real world. So, on these assumptions, *flat* satisfies the all-or-nothing condition. The same is true for the complementary predicate to *flat*, the minimum criterion predicate *uneven*: For a surface x to be uneven it is necessary and sufficient that x have some unevenness, however slight. But surely every physical surface will have *some* unevenness somewhere, so every physical surface will be uneven: the positive extension of *uneven* consists of its entire application domain.

In light of the above discussion, it might have been thought that at least for those absolute adjectives for which the all-or-nothing condition is fulfilled, we get a resolution of the sorites paradox along the lines sketched above. But somewhat ironically, absolute adjectives are predicates for which such a resolution of the paradox is neither plausible nor necessary (cf., Kennedy 2007). Imagine (*per impossibile* perhaps) a perfectly flat surface x_0, a severely uneven surface x_N, and a series of surfaces $x_1,\ldots,x_{N-1}$ connecting them, where the differences between the successive elements are as small as you like. Here the second, inductive premise lacks plausibility, since there is one instance of it – the conditional that if x_0 is flat, then x_1 is flat – for which there is no clear intuitive backing. So, since in this case the second premise seems plainly false, there is nothing paradoxical about its false conclusion.

Notice, however, that if absolute adjectives were to satisfy the all-or-nothing condition in the real world, then these adjectives would have no meaningful application in practice. If everything is dirty, why make a point of saying that something is dirty, and if nothing is ever truly clean, why would competent speakers engage so often in stating falsehoods? In practice, the maximum and minimum criteria that govern absolute adjectives are applied leniently or sloppily. How lenient will depend on context. And here too what depends on context is some kind of tolerance: Context will determine, to a greater or lesser extent, how much dirt we tolerate – e.g., on the shirt or on the carpet – before we stop calling the thing clean and get ready to start calling it dirty.

When absolute adjectives are used in this lenient way, then of course they will be vulnerable to Sorites effects again. If the difference between clean and dirty is just a matter of how much dirt, then *clean* and *dirty* turn into tolerance-governed predicates in the earlier sense that small differences make no difference to their satisfaction conditions.

2.2 Denying Inductive and Conditional Premises: Partial Knowledge or Partial Information?

A second strategy for dealing with the Sorites paradox is to acknowledge the base premise of (12) and (13) as true and the conclusion as false, but to argue that the other presmises (the inductive premise in (12) and one or more of the conditional premises in (13)) are false. A number of approaches follow this strategy. Salient among them are two: epistemicism and Supervaluationism.

2.2.1 Epistemisicm.

We already saw that for epistemicism the Sorites isn't a problem. There are no vague predicates. So the second premise of (12) is false, and so is one of the conditional premises in (13). It is just that it is fundamentally impossible for us to know which of those conditionals is false. And likewise, we cannot know which instance of the second premise of (1) is responsible for its falsehood. Another example is (14a), which expresses that being one day apart in age is within the tolerance relation for the predicate *child*. On the epistemicist view (14a) must be false as well. Moreover, and for the same reason, (15), the negation of (14a), is true, although we cannot hope to find any particular n which instantiates this existential statement.

(14) a. For any n, if an n days old human being is a child, so is an n+1 days old human being.
b. If a 1 day old human being is a child, then that human being is also a child when it is 2 days old.
c. If a 1 day old human being is a child, then that human being is also a child when it is 36,500 days old.

(15) There is a number *n* such that an *n* day old human being is a child but is no longer a child at *n* + 1 days.

Epistemicists (Williamson 1994, 2000; Graff-Fara 2000; Sorensen 1988, 2001) accept this last consequence. But it is in particular the inescapability of this consequence of Epistemicism that epitomises what its critics find hard to swallow.

2.2.2 Supervaluationism.

As opposed to Epistemicism, Supervaluationism assumes that vague predicates are genuinely partial (Lewis 1970; Fine 1975; Kamp 1975; Kamp and Partee 1995; Keefe 2000). This assumption is implemented using partial models in which a vague predicate P has a positive extension $[\![P]\!]^+$, a negative extension $[\![P]\!]^-$, and a truth value gap $[\![P]\!]^?$. For simplicity we assume that the application domain of P includes all of the Universe U of the model. In the formal developments that follow, this will from now on be taken for granted. Distinctive of the Supervaluation approach is that it considers, in conjunction with each such partial model M_0 a set *S* of models that extend M_0 (where a model M' is an *extension of* a model M iff for each predicate P, $[\![P]\!]^+_M \subseteq [\![P]\!]^+_{M'}$ and $[\![P]\!]^-_M \subseteq [\![P]\!]^-_{M'}$).[6] We follow the tradition in confining ourselves

[6] See for example in Stalnaker (1975, 1978); van Fraassen (1969); Kamp (1975); Fine (1975); van Benthem (1982); Veltman (1984, 1996); Groenendijk and Stokhof (1984); Landman (1991), Barker (2002), and others.

to the special case of a language L of first order predicate logic with one or a few partial predicates. Unless mentioned otherwise, we assume that L has just one partial predicate P.

The set *S* of extensions can be thought of in various ways. Stalnaker (1975) shows how a construction of this kind can be used to model many aspects of information states in a background context c. In this interpretation the members of *S* are the states of affairs, or 'possible worlds', that are compatible with the given information state. When we think about information states as states of information about the partial predicate P, then the models in *S* should be thought of as giving the different ways in which the truth value gap of P in the 'ground model' M_0 could be resolved in a manner that is consistent with that information state. A different way to think of *S* is as the set of possible modifications of the semantics of the predicate P. This could be either an enduring change of the language L to which P belongs – that is, L is changed into a new language L' – or as a more 'local' effect: as the modification of a given context c of use for P to a more specific context c', in which certain questions about P's extension that remained undecided in c have now been settled. Such special contexts will typically last only for the duration of the given utterance or conversation and reverse to the neutral context reflected by the ground model when the circumstances of the utterance no longer prevail.

This latter way of thinking about *S* suggests that we should distinguish between contexts of use and the information that contexts carry about the extensions of partial predicates: In general this latter information is only one part of the total information carried by the context. We capture this distinction by assuming that each context c determines a partial model Mod(c), which incorporates the information in c about the extensions of the predicates of L. We write c ≤ c' to convey that c' carries at least as much information as c; 'c ≤ c' entails that Mod(c) ≤ Mod(c'), where for two partial models M and M' for L, M ≤ M' iff for every predicate P of L $[\![P]\!]^+_M \subseteq [\![P]\!]^+_{M'}$ and $[\![P]\!]^-_M \subseteq [\![P]\!]^-_{M'}$. From now on we assume that the contexts form a set C and that S_C is the set of corresponding models: S_C = {Mod(c): c ∈ C}. When M ≤ M', we say that M' is an *extension of* M.

These stipulations lead us to structures of the form <C,≤,Mod>, with ≤ a partial order of C and Mod a function from C to partial models of L. We assume in addition that C always has a ≤-minimal context c_0, which stands in the relation ≤ to all members of C. We refer to Mod(c_0) as the *ground model* of <C,≤,Mod>.[7]

The models in S_C can be either partial or complete (that is, classical) models. We refer to the set of complete models in S_C as T_C. The members M of T_C are called *precisifications* (*of* those models Mod(c) in S_C such that Mod(c) ≤ M). With these additional notions we can represent our structures as 5-tuples <C,≤,c_0,Mod,*T*>.

In the application of these models to the analysis of vagueness, the relation ≤ between models should **only** reflect the resolution of the truth value gaps of partial predicates. That is, when M ≤ M', then the universes of M and M' should be identical. In fact, it is useful to restrict attention to Supervaluationist models <C,≤,c_0,Mod,*T*> in which all models have the same universe. In other words, each model Mod(c) in S_C has the same universe as Mod(c_0). Models that satisfy this additional constraint of model constancy will be called *Supervaluationist models*, or *SV models*.

[7] Sometimes the assumption is made that Mod(c_0) is minimal in an absolute sense: for every partial predicate P, $[\![P]\!]^+_{Mod(c_0)} = [\![P]\!]^-_{Mod(c_0)} = \emptyset$; however, we do not adopt this as a general assumption.

The semantics of vague predicates can depend on context in a variety of ways, and not only on the extensions that the contexts assign to them via their associated models. That is why it is important to be able to distinguish in general between contexts and models. But in many applications of the Supervaluation methods to vagueness it is only the extensions that matter. In those applications, it is possible and convenient to identify contexts with their associated models. This identification leads to SV models of the form $M^* = \langle C,\leq,c_0,T\rangle$, where C is now a set of partial models and ≤ the model extension relation. A further specialization is to SV models in which all models in C other than c_0 are complete. Such SV models are often represented as $\langle M_0,S\rangle$, where *S* is a set of complete extensions of M_0.[8]

There are two distinct uses to which SV models can be put. The first is to define notions of *supertruth* and *superfalsity*: a sentence φ of L is *supertrue in* the context c of an SV model M^* iff φ is true in all members c' of T such that c ≤ c'; and φ is *superfalse in* c iff φ is false in all members c' of T. Furthermore, φ is said to be *supertrue/superfalse in* M^* iff φ is supertrue/superfalse in c_0.

For simple languages like our language L of first order predicate logic truth and falsity in partial models are *monotone*: if c ≤ c' and φ is true/false in c, then φ is true/false in c'.[9] Evidently, when truth and falsity are monotone, then for each c in C and sentence φ, if φ is true/false in c, then φ is supertrue/superfalse in c. But the converse does not hold in general. For instance, every sentence of L that has the form of a tautology will be supertrue in any SV model M* (in every c of M*) and every sentence that has the form of a classical contradiction will be superfalse. In particular, $P(a) \vee \neg P(a)$ will be supertrue and $P(a) \,\&\, \neg P(a)$ superfalse in all SV models M*, even when P is vague and *a* is in the truth value gap of P in the ground model. Note that supertruth and superfalsity lack many of the compositionality properties of the classical notion of truth for first order logic. For instance, the supertruth of $P(a) \vee \neg P(a)$ does not entail the supertruth of one of its disjuncts, whereas the truth of a disjunction p ∨ q entails that either p or q must be true. In a typical SV model in which *a* is in the truth value gap of P in the ground model, there will be t in *T* where *a* is in P's positive extension and also t in *T* where *a* is in P's

---

8 In such applications, the 'intermediate models' of SV models, i.e., those $c \in C$ such that $c_0 < c < t$ for some $t \in T$, do not do any work and can be ignored. The result is a structure $\langle M_0,S\rangle$ of the kind just described.

9 Truth for L in partial models M requires a 'two-sided' recursion, with separate clauses for truth and falsity. For instance, the clauses for (i) atomic formulas of the form P(τ), (ii) negation, (iii) conjunction and (iv) existential quantification must be stated as follows. (Throughout, g is an assignment of entities in the universe of M to the variables of L, g(τ) is the value g assigns to the variable x if τ = x, and the denotation of the indivdual constant c if τ = c and g[d/x] is the assignment that is like g except that g[d/x](x) = d.)

(i) $[\![P(\tau)]\!]_{M,g} = 1$ iff $g(\tau) \in [\![P]\!]^{+}_{M}$

$[\![P(\tau)]\!]_{M,g} = 0$ iff $g(\tau) \in [\![P]\!]^{-}_{M}$

(ii) $[\![\neg\varphi]\!]_{M,g} = 1$ iff $[\![\varphi]\!]_{M,g} = 0$

$[\![\neg\varphi]\!]^{M,g} = 0$ iff $[\![\varphi]\!]_{M,g} = 1$

(iii) $[\![\varphi \,\&\, \psi]\!]_{M,g} = 1$ iff $[\![\varphi]\!]_{M,g} = 1$ and $[\![\psi]\!]_{M,g} = 1$

$[\![\varphi \,\&\, \psi]\!]_{M,g} = 0$ iff $[\![\varphi]\!]_{M,g} = 0$ or $[\![\psi]\!]_{M,g} = 0$

(iv) $[\![(\exists x)\varphi]\!]_{M,g} = 1$ iff there is some d in the universe U_M of M such that $[\![\varphi]\!]_{M,g[d/x]} = 1$

$[\![(\exists x)\varphi]\!]_{M,g} = 0$ iff for all d in U_M, $[\![\varphi]\!]_{M,g[d/x]} = 0$.

These truth clauses are characteristic of historically earlier systems, including strong Kleene (K3) and logic of paradox (LP), the relations to which are further discussed shortly, as well as to the 4-valued Dunn-Belnap semantics (Benlap 1977).

negative extension. A similar deviation from what holds for truth in the absence of vagueness is found in its interaction with quantifiers, e.g., an existential sentence can be supertrue without any of the corresponding instances being supertrue.

The notion of supertruth can be used to define logical validity:

(16) B is a *logical consequence of* Γ iff for every Supervaluationist vagueness model M*, if every premise A in Γ is supertrue in M*, then so is B.

The consequence relation defined in (16) turns out to be that of classical logic. This is one possible way of defending classical logic in the face of vagueness, but it is not the only one. A second definition will be given in (17) below.

Some of the criticisms of the Supervaluation approach have targeted the notion of supertruth[10]. Supertruth, it is contended, isn't really a notion of truth at all, which could serve as a proper explication of the informal notion of truth as we know and use it and that philosophy and logic should analyse and clarify. The factors that disqualify supertruth as a notion of truth include both the way supertruth is defined, involving a universal quantification over virtual models, which represent possible interpretations that *could* be assigned to the predicates in question, and certain formal properties of supertruth that are incompatible with those that any intuitively plausible notion of truth should have.

But is this to be seen as a criticism of the Supervaluation method as such? Not so, we think. Supervaluation and Supertruth should not be conflated. It is important in this connection to note that the Supervaluation approach can be used to characterise logical validity without making use of supertruth. This is the second way referred to above to make use of the same SV models M* = $\langle C, \leq, c_0, T \rangle$ that are employed in the first use. But we now define logical validity directly on the complete extensions of such models:

(17) The argument Γ |= B is *supervalid* iff for every SV model M* and every t in T_{M^*}, if every A in Γ is true in t, then so is B.

This definition is like the one in (16) in that it too gives us classical logic. But it does this in a somewhat different way. Technically, it is more 'local' (cf., Williamson 1994, Varzi 2007, and Cobreros 2011). Rather than requiring, in the way of (16), the preservation of supertruth, which itself involves the idea of quantification over the complete models of particular SV models M*, it requires preservation of (classical) truth in each complete model of every SV model. But apart from formally circumventing the notion of supertruth, the definition also has an intuitive motivation that is not easy to recover from the definition in (16).

Logical deduction is a matter of form. We can engage in deduction even when we have no idea what some of the content words in the premises and conclusion of an argument mean; and *a fortiori* we can do so without knowing all there might be to know about those words. In this sense, most of deductive reasoning we engage in has a hypotheticaI dimension: We do not know the exact truth conditions of either premises or conclusion. But we do know that these truth conditions must be related in such a way that if the premises come out as true, given the truth

[10] Sorensen (2013) argues that when Supertruth is accepted as the relevant notion of truth in the presence of vagueness, then the Supervaluation account either commits us to the truth of the proposition that a one year old is not a child or else to the truth that a seventy year old person is a child. Neither is acceptable. So supertruth ins't a good explication of truth.

conditions they prove to have and given what the world may turn out to be like, then the corresponding truth conditions for the conclusion will guarantee its truth in these same circumstances. In particular, it is quite immaterial in this connection whether the information we may lack about the meanings or extensions of certain words occurring in an argument is due to our own ignorance or to the fact that there just isn't any more information to be had than what we already have. All that we need to assume in order to justify the deductive inferences we draw is that the accommodations needed to obtain determinate truth conditions for the premises are the same as those needed to obtain determinate truth conditions for the conclusion.

Within a Supervaluationist setting it is also possible to define logical validity in yet other ways. (18) is one possible definition of an *internal* logic, which is based directly on the ground models of SV models and makes no reference to any extensions of it. So (18) isn't really Supervaluationalist in spirit. But SV models suffice as its inputs.

(18) $\Gamma \models B$ is *internally valid* iff for every partial model M, if every A in Γ is true in M, then so is B.

This definition imposes stricter constraints on validity than supervalidity and thus defines a weaker logic, which is known as 'Strong Kleene' (Veltman 1984).[11] (18) is only one way in which the internal logic of partial models can be plausibly defined.[12] For someone who wants to advocate such an internal logic as the 'true' logic of vague predication this presents a conceptual problem: How is he to argue for a motivated choice between these different options?

Is there a way to argue that definition (17) is conceptually superior to definition (18) or any of its variants? For a language of predicate logic with vague predicates the argument we gave in favour of (17) seems persuasive to us. But the issues change when we consider languages which include operators like Def. In this chapter we do not go into the semantics and logic of such operators just as we do not go into higher order vagueness. But it should nevertheless be pointed out, we feel, that the introduction of operators like Def marks an important shift in both the technical and the conceptual aspects of the logic of vagueness.

The distinction between the Supervaluation approach (i) as a way of reanalysing the concept of truth as supertruth and then defining logical consequence in terms of it and (ii) as a method for analysing logical validity without making use of supertruth is important also for the account that Supervaluationism is able to give of the Sorites paradox. The standard Supervaluation account, already mentioned above, is that some premise or premises other than the Base Premise is or are, contrary to appearances, false. For an argument of the form (12), this is the Inductive Premise. But for arguments of the form (13), the account cannot be quite formulated in this exact form. To

[11] One salient difference between classical logic and Strong Kleene is that no classical tautology is valid in the Strong Kleene logic (in the sense of following logically from the empty set of premises).

[12] One alternative definition is given in (i):

(i) $\Gamma \models B$ is internally valid iff for every partial model M in which B is false, at least one of the premises in Γ is false.

The logic generated by this definition is different from the Strong Kleene logic generated by (18). It is known as 'LP'. LP differs from Strong Kleene, and is in this respect more like classical logic, in that it verifies all classical tautologies; but like Strong Kleene, it is a good deal weaker than classical logic. There are still other ways to define internal validity besides these two, which yield yet other logics. We refer to the literature for details (see in particular, Varzi 2007 and Asher et al 2010).

see this, consider a SV model M*, whose universe is the set of the positive natural numbers and that P represents the predicate 'is a number n such that n grains of sand can make a heap'. In keeping with our intuition that the Base Premise of (12) is definitely true, and assuming that its conclusion is definitely false, the positive extension of P in the ground model c_0 of M* should be some segment [k,∞) that is neither too small nor too big (e.g. k is somewhere between 50 and 1,000,000), and the negative extension of P in the ground model should be some interval [1,l], where $1 < l < k$. Furthermore, for each t in *T*, there is a number r in the interval [l,k) such that the negative extension of P in t is the interval [1,r] and the positive extension [r+1, ∞). Conversely, we assume that for each r ∈ [l,k) there is a t ∈ *T* in which P has the mentioned extensions. It is clear that in the ground model of M*, P(*1,000,000*) is true and P(*1*) false. Furthermore, the Inductive Premise of (12) is superfalse in M*. For let t be any precisification in *T*. Then for some r ∈ [l,k), P(*r+1*) is true in t and P(*r*) is false. So the conditional P(*r+1*) → P(*r*) is false in t and, therefore, also the Inductive Premise, which is the universal quantification over such conditionals. So the Inductive Premise is false in every t ∈ *T*; that is, the premise is (super)false in M*. So M* isn't a counterexample to the validity of (12). And since we may expect that any model which describes a Sorites situation yields this same conclusion, there is no reason here to doubt the validity of the argument.

We can argue in much the same way that arguments of the form (13) do not present countermodels to validity either. But to show this we have to proceed a little differently. This is because in a model like M* none of the premises P(*r+1*) → P(*r*) will be superfalse. But even if no one of these premises is superfalse by itself, the conjunction will be superfalse. Or, put only slightly differently, the set of these premises cannot be jointly true: there is no single complete model t in M* such that these premises are true in t. So that a model like M* isn't a countermodel to (13) follows from the following principle:

(19) A false conclusion can be validly deduced from a set of premises that cannot be jointly true.

(19) can also be used to explain why models like M* aren't countermodels to (12). In (12), the Inductive premise does the work that the conditional premises P(*r+1*) → P(*r*) do in (13). So (19) can be applied to the case of (12) by taking the set of premises that cannot be jointly true to be the singleton set of which the Inductive premise is the only member. This set 'cannot be jointly true' in M* because its one member is false in all members t of T. We can formulate this by saying that the Inductive premise is superfalse in M*, but we think enough has been said to make it clear that the solution that the Supervaluation method offers for the Sorites not only can do without the notions of supertruth and superfalsity, but that in accounting for (13) these notions do not even deliver what is really needed. Here too we see that the Supervaluation method is not wedded to the concept of Supertruth as the incarnation of truth in the presence of vagueness.

These last considerations were aimed to show that the Supervaluationist resolutions of the Sorites paradox are coherent. The arguments are valid, but in models that capture our intuitions about the conditions in which their premises and conclusions get their truth values, the premises can never be all true. But arguably coherence isn't all that we want. Perhaps the strongest objection against the solution that Supervaluationism offers of the Sorites paradox is that it fails to do justice to the strength of our sense that the conditional premises of arguments of the form (13), and perhaps also the Inductive Premise of (12), are true. Consider once more one of the conditional premises in (13):

(20) If n+1 grains of sand can make a heap, then n grains can make a heap too.

The intuition that such a sentence is true is so strong, so the objection goes, because there is a dynamic aspect to its meaning: acceptance of the antecedent of (20) in a context c produces a context change from c to a new context c’ in which n has become established as a member of the positive extension of the relevant predicate P. When, in this new context c’, the question is raised whether P is true of some other number n’ that stands in the tolerance relation to n, the answer can only be affirmative. Anything within the tolerance range of something that is an established member of the positive extension of P must also be classified as P. Dynamic approaches to the Sorites paradox are among the topics of Section 2.2.4.

2.2.3 Subvaluationism

An approach to vagueness that has many formal features in common with Supervaluationism is **Subvaluationism**; the name has been chosen to emphasise these formal similarities. However, the informal motivation behind Subvaluationism is very different from that behind Supervaluationism. Conceptually, Subvaluationism is embedded within **Paraconsistency**, a general view of logic in the face of paradox (Priest 1987). In paraconsistent logic, paradoxical combinations of sentences can be represented and adopted as premises without the consequence that everything becomes derivable.[13]

As regards vagueness, the core intuition of the paraconsistent approach is that when *a* is a borderline case of P, then P(*a*) is *both* true and false (rather than neither true nor false, as in the Supervaluationist approach). One way to make this intuition explicit is something we could have done in our presentation of Supervaluationism too, but which doesn’t seem quite as natural there: introduce a third ‘truth value’, in addition to the truth values we already have – 1 for true (and not false) and 0 for false (and not true) – in order to cover such cases of joint truth and falsity. The new truth value can be denoted as ‘(1,0)’, so that the *truth value space* becomes {1, (1,0), 0}. Three-valued models based on this truth value space for the language L described in section 2.2.2 will assign to predicates functions that map their application domain into {1, (1,0), 0}. The borderline cases of P in such a model will now be those individuals which the function for P maps to (1,0). A three-valued truth definition, which closely resembles that for partial models given in footnote 9, then extends these truth value assignments to all sentences of L.

It is possible to associate with such three-valued models sets of complete extensions in much the same way as that is done in the Supervaluationalist approach. Going from M to one of its complete extensions now means that each entity *a* such that P(*a*) has the value (1,0) in M gets either the value 1 or the value 0 in the extension. Evidently the structures thus obtained are isomorphic to the vagueness models <M,*S*> of the Supervaluation approach. So exactly the same concepts can be defined from them. But what concepts will seem useful will depend on the underlying philosophy, and in view of what it means for the Subvaluationist that a sentence has the value (1,0), it seems natural to introduce notions of ‘subtruth’ and ‘subfalsity’as follows.

(21) Let M** be a pair <M,*S*> where M is a three-valued model and S is a set of complete extensions of M. A sentence φ of L is *subtrue in* M** iff there is some model M’ in *S* such that φ is true in M’. (And likewise for subfalsity.)

[13] An early formulation of formal logic with this property, known as LP, can be found in Priest (1979.)

While the Supervaluationist approach has put a strong emphasis on the 'outer logic' which can be defined using the complete extensions of SV models, the Subvaluationists have shown a preference for an internal logic, which can be defined directly on the three-valued models themselves (see Hyde 2008, who endorses and later rejects an outer logic). The preference has been for the logic LP (see footnote 12), which has an independent motivation as a general logic of paradox. One of the uses of LP is in an account of the Sorites paradox. As this too is an account that advocates a change of logic, we wait with the details until Section 2.3, which is devoted to Sorites responses of this type.[14]

2.2.4 Contextualism:

By Contextualism we here understand those approaches towards the Sorites problems that rely on the context dependence of vague predicates. We already noted the possibility of interpreting conditionals like $P(a_{i+1}) \rightarrow P(a_i)$ as true because their consequent is true in the context established by their antecedent. But this is only one way in which context has been argued to play a part in Sorites phenomena. Some of the various different ways in which vague predicates can depend on the context of use have no direct bearing on the Sorites paradox, but they nevertheless deserve mention and will become relevant in Section 3.

1. The type of context dependence of vague predicates that has been prominent in the formal literature on vagueness is most clearly visible in prenominal occurrences of adjectives, as in *small elephant, tall basketball player,* and *hot bath*. Kamp (1975) and Siegel (1976, 1979) are early proposals for a treatment of prenominal adjectives as vague context-dependent predicates whose contexts are determined wholly or largely by the nouns that follow them. The semantics of such combinations AN involves an adaptation of the extension of the adjectival predicate A to the extension of the noun N.

1a. More generally, the satisfaction conditions of an adjective A (when used in its positive as opposed to its comparative or superlative form) depend on a *comparison class* C, within which the adjective should make a meaningful distinction in the sense that both the intersection of C with the set of things that satisfy A in C and the difference between C and that intersection are significant portions of C.[15] The case where C is the extension of the noun following a prenominal adjective is a special but frequent instance of this.

1b. In some adjective noun combinations, such as *talented pianist* or *skillful cobbler*, the noun plays a somewhat different role than the one indicated under 1a. A skillful cobbler is someone who is skillful *at* cobbling. You can be a cobbler and a dart player, but a skillful cobbler while not a skillful dart player. What is involved are different skills, not a single skill of which you have enough to qualify as a skillful cobbler, but not as a skillful dart player (Lewis 1969).

1c. A related phenomenon can be observed with 'multi-dimensional' adjectives like *clever*. People, animals and actions can be called clever for varying reasons, reaching from prowess at mathematics to social adroitness (possibly including a talent for manipulating or misleading

[14] We note that not all who endorse LP as a logic of vagueness are Subvaluationists, just as one does not have to be a Supervaluationist to endorse the Strong Kleene logic. For a different setting in which these logics arise see Cobreros et al., (2010).

[15] See Klein (1980) and van Benthem (1982) for procedures that compute such a 'practically meaningful' extension of A in the context of a comparison class C from C and the meaning of A.

others). Often the context will make clear what kind of cleverness is intended; and sometimes what is intended is some kind of mixture of the different criteria.

2. When a vague predicate is governed by tolerance, how much tolerance is it governed by? As noted in Section 1, that too varies with context. The size of the tolerance relation (i.e. how far two elements that stand in the tolerance relation can be apart from each other) is often referred to as the **granularity** of the predicate. For instance, a context may carry the commitment to count two human beings *a* and *b* as either both old or both not old when their ages are a year or less apart. In this case, the contextual tolerance relation spans intervals of a year or more. But in other contexts the tolerance relation may be a narrower one, which carries a correspondingly lesser commitment. For instance, the commitment to count *a* and *b* as both old or both not old if their ages differ by a month or less (or by a day or less), is not as much of a commitment that it is to count them as both old or both not old whenever their age difference is less than a year.

3. Another form of granularity arises with predicates that are derived from terms that denote elements of space or time, such as *local* or *noonish* (Sorensen 2001; Ripley 2013). How local must something be to count as an instance of *local*? how close to 12.00 a.m. must a time be to count as noonish? This too varies with context. But in these cases the granularity takes the form of narrower or wider concentric circles around some spatial dot or temporal point. Granularity of this kind is found not only in adjectives like *local* or *noonish*, but also in the terms from which they are derived. 'He came at noon' can count as true even when the exact arrival time was 11.57, or 12.03.[16]

All of these context dependence effects are different from the dynamic effects touched upon in the closing paragraphs on Supervaluationism in section 2.2. They are denied direct relevance to the Sorites paradox in Shapiro (2005), where the distinction is drawn between 'external' and 'internal' contexts. The external context, including the contextual factors in 1.-3., is typically fixed in a situation in which the Sorites arises and remains unchanged in the course of it. The Sorites problems arise because of the internal context, which is responsible for the dynamic aspects of vagueness (but see the discussion in section 2.2.4.2 whether it is really true that the factors in 1-3 have no direct bearing on the Sorites).

Actual situations in which the Sorites paradox can become an acute problem for participants in a discourse are referred to as **forced marches** (Horgan 1994). A forced march arises when it is beyond question that an actual entity a_1 satisfies a predicate P; it is equally beyond question that an actual entity a_N does not; and for i = 1,…,N-1, the actual entities a_i and a_{i+1} stand in the tolerance relation for P. It is obvious that our world is densely populated with such configurations (given particular tolerance relations for the tolerance-governed predicates). How much of a difficulty these situations actually do present to speakers who find themselves in such a situation and want to talk about it is itself a topic of theoretical debate. But before anything can be said about this debate we must first make more precise what the problem is supposed to be.

In the vagueness literature, the forced marches discussed have been mostly situations brought about in psychological experiments. Prominent among experiments that have actually been

[16] The granularity associated with such terms depends in large part on whether the term denotes a 'round number' (in relation to some convention for measuring the magnitude in question, such as, in the case of *noon*, time). For instance in a context in which an arrival at 12.03 could be described as 'He came at noon', the description 'He came at 12.01' would not be acceptable (Lasersohn 1999; Krifka 2002). The round number effect also arises for terms denoting the natural numbers. 'A hundred people came to the lecture' may be acceptable when the exact number was 98, but 'Ninety eight people came to the lecture' wouldn't. As Lasersohn (1999) puts it, some terms have a 'halo' around what is to be regarded as their denotation in the strict sense (see section 3.2).

performed are those involving colour predicates (cf., Egre, Gardelle and Ripley 2013). Here is one specific way in which such an experiment can be set up. The material consists of a series of N coloured squares running from an unequivocally red square a_1 to an unequivocally orange square a_N, while successive squares a_i and a_{i+1} look exactly the same colour when placed next to each other. The experiment then proceeds by first showing square a_1 to the subject, asking him to say whether or not it is red. This will invariably elicit an affirmative answer. One then puts a_2 beside a_1, asking the same question about a_2. Again the answer will normally be affirmative. Then a_1 is removed and a_3 put next to a_2 and the same question is repeated. In this way the experiment is continued until the point where the subject expresses a doubt that the square a_i is red or even denies this. At this point the experimenter may halt the experiment, noting at what square a change in response occurs; or he may continue by probing the subject's reasons for his change of heart, and perhaps return to the judgement he had delivered about the preceding square a_{i-1}. Other continuations are possible as well, and there are also different ways of conducting the experiment, for instance, by not removing squares, but rather displaying longer and longer sequences of squares $<a_1,.., a_i>$ (Raffman 1994, 2005, 2014). But these variations do not matter for what we want to say here.

Important for what we want to say next are two points. First, if we make the familiar and plausible assumption that any two colour patches whose colour shades cannot be told apart (by subjects with normal vision in normal viewing conditions) stand in the tolerance relation for ordinary colour predicates like *red* and *orange*, and if it is strictly true that these colour terms are governed by tolerance, then the judgement changes to which forced march experiments lead more or less without exception are violations of the semantics of those predicates. Either the subjects of forced march experiments violate the semantic rules of their language or else it is part of the semantics of tolerance-governed predicates that the tolerance principles governing them are not absolute, but capable of being overruled by other principles with a higher priority.

Support for this latter position can be found in the work of Sainsbury (Sainsbury 1990). Sainsbury notes that there are situations in which it seems fully legitimate to adopt a sharp cut-off point for colour terms and other vague predicates of which there can be no doubt that they are sensitive to tolerance. An example is that of the owner of a paint shop who decides to store the reds on the top shelf and the oranges on the shelf below. This requires a decision, arbitrary to at least some extent, as to where the cut between the reds and the oranges should be made. Once that decision has been taken it is better to stick to it for as long as one can, at least when it comes to restocking the shop. But that doesn't diminish the freedom of choice that exists *at* the time when the decision is made. In last analysis there is nothing in the semantics of *red* and *orange* that prohibits such decisions. And if that is so in the situation of the paint shop, why should it be so radically different in the situation of a forced march, where the price for sticking to the tolerance principle *à tor et à travers* is bound to get you into serious trouble?

What is it that countermands the power of tolerance in forced march situations? This is a question on which, we believe, no detailed answers currently exist. Certainly there isn't wide agreement on some particular answer (see Egre, Gardelle and Ripley 2013 for actual experimentation and discussion.) Coming up with a convincing answer appears to us as one of the important tasks for the theory of vagueness and tolerance.[17]

[17] Raffman (1994) diagnoses the judgement changes that occur in the course of forced marches as a species of Gestalt switch, akin to what can be observed when subjects are made to stare at a Necker cube or certain other ambiguous images. On this view the judgement switch occurs when the new pair (a_i, a_{i+1}) is submitted to the

A second point about forced marches is that the linguistic resources involved can be extremely simple. Typically, in a forced march experiment the only sentences ever uttered are those occurring in question-answer pairs like: A: ‘Is this red?’; B: ‘Yes’/’No’/’I am not sure.’ These are just atomic predications. None of the logical constructs – conjunction, conditionals, universal quantification – that occur in the sorites arguments on which so much of the sorties literature has focused, play any part here. When this point is appreciated, the challenges that are presented by Sorites arguments in their traditional forms, as in (12) and (13), take on a new complexion. In these arguments, logical constructs play a central part, for it is they that permit, or seem to permit, the deduction of the conclusions from the premises. So the focus has to be on the way these constructs behave in the presence of vagueness. The Supervaluationist response meets the challenge by modeling vagueness in such a way that the classical semantics of the constructs can be preserved and resolves the paradox by arguing that the premises cannot be jointly true. But if one is not prepared to place the blame as squarely as that on the truth values of the premises, then it is no longer clear whether the classical semantics of the logical operators can be upheld. Contextualist approaches to the Sorites differ on this point, as we will see next.

2.2.4.1 Dynamic Contextualist Accounts

As we see it, the dominant trend within contextualism is dynamic, and the core of dynamic accounts is what we just surmised about the mechanisms involved in forced march conversations: the extensions of vague predicates are in general only partially determined, but the extent to which they are determined varies with the context. More specifically, the participants in a conversation can agree to treat entities mentioned in the conversation, or those newly brought into it, as positive or negative instances of P, thereby changing the context into one of which this commitment is part. If, moreover, P is governed by tolerance, then such contextual changes may have a secondary effect: for an entity *b* with regard to which it was previously undecidable whether it should be counted as P, it may now be possible to answer this question, because the new context carries the commitment that some individual *a* that stands in the tolerance relation to *b* belongs to, say, the positive extension of P; in that case, assuming that the tolerance principle is in force, P(*b*) should be regarded as true. But that means that if in the given context someone raises the question whether P(*b*), and the participants act according to the dictates of tolerance and jointly accept the truth of P(*b*), then that will lead to yet another context, in which there is a commitment that *b* now also belongs to the positive extension of P.

In this way, the interaction between context-dependent truth judgements and the context-modifying commitments that arise when these judgements are expressed and acknowledged as

---

subject’s judgement after a_{i-1} has been removed. Once the Gestalt switch takes place, *neither* member of the pair (a_i, a_{i+1}) will look unequivocally red to the subject any longer; and when the experimenter then returns to the pair (a_{i-1}, a_i), then the Gestalt switch that has just occurred will tend to have the effect that now neither element of this pair, a_{i-1} no more than a_i, seems to support an unequivoval judgement of ‘red’.

Of course, even if it is true that Gestalt switches lead to the changes in the judgements subjects actually express, this does not prove that the *semantics* of *red* and other colour predicates is compatible with such changes. There might be a fundamental rift between our understanding of how such predicates should work and what is dictated by our colour experiences in forced march situations.

Cobreros et al. (2012) advocate tolerance as a soft and defeasible constraint, and Raffman (2014) even considers the necessary category-switch in forced marched sorites as a primitive datum for an adequate theory of vagueness. She would, therfore, rather ask what makes the power of tolerance so appealing in Soritical contexts, despite its being doomed to be given up.

correct will propel contexts and judgements onwards, and to eventual disaster if nothing puts a stop to it. Exactly what the stopping mechanism is, what triggers it and what this entails for the semantics of tolerance-governed predicates is part of the open question mentioned above.[18]

As described so far, these contextualist accounts say next to nothing about the semantics of the language L. All they specify are partial truth and falsity conditions for atomic formulas of the form P(a). In fact, there are various directions in which such an account can be expanded to one that covers more substantive languages or language fragments, such as the logically complex sentences of our language L. Here is a direction one might consider taking.[19] Suppose that c_i are the successive contexts that, among other things, determine the positive and negative extension of the vague predicate P in the course of a conversation. (We assume that some account is given of how conversational moves determine the change from c_i to c_{i+1}.) Each c_i is assumed to determine the positive and negative extension of P in it. So each c_i determines a partial model M_i in which P has a positive and a negative extension. We can now associate with each of these M_i's an SV model M^*_i and use these models in whatever way we see fit to account for the semantics and logic of L. In this way, we can, if we want to, retain classical logic. As the conversation proceeds, the ground models of the SV models become increasingly determinate, so the set of atomic and non-atomic sentences of L that have a definite truth value grow. And if the conversation takes the turn of a forced march, then this would lead in the end to untenable models M_i if it weren't for some overriding principle to halt the process.

In this way the classical definitions of the logical operators are preserved (just as they are in the Supvaluationist approach). But a number of contextualist accounts of sorites arguments are more radically dynamic in that they also argue for a dynamic reinterpretation of logical operators. Kamp (1981) and Shapiro (2005) share a sense of the need for a dynamic revision of the conditional along the lines of the following truth and falsity clauses:

(22) a. $[\![\varphi \rightarrow \psi]\!]_{M^*,c} = 1$ iff either $[\![\varphi]\!]_{M^*,c} = 0$ or $[\![\psi]\!]_{M^*,c+[\![\varphi]\!]} = 1$
b. $[\![\varphi \rightarrow \psi]\!]_{M^*,c} = 0$ iff $[\![\varphi]\!]_{M^*,c} = 1$ and $[\![\psi]\!]_{M^*, c+[\![\varphi]\!]} = 0$

[18] Soames (2002) describes the dependence of the extensions of vague predicates on context as a form of indexicality. This has been criticised by Stanley (2003) as untenable because soritical situations can be described also in a form like that in (i).

(i) If a_1 is a heap then a_2 is too, and if a_2 is, then a_3 is, and if a_3 is, then a_4 is, … and then a_N is.

It is a feature of indexical expressions (such as first or second person pronouns and deictic demonstratives) that they cannot change reference between two occurrences one of which is overt and the other part of the tacit material that is missing from ellipsis constructions like VP Deletion or Gapping and that needs to be recovered to interpret those constructions. So if *heap* was indexical in the sense that its extensions were indexically determined, then there could be no relevant change in context between the antecedent of the first conditional in (i) and all the following clauses (each of which is elliptical). Clearly, the context-dependence at issue does not follow the principles that, essentially following (Kaplan 1989) are assumed for indexicals. In fact, other contextualist proposals for dealing with vagueness do not make the assumption that the context dependence of vague predicates is a form of Kaplanian indexicality.

[19] We are not sure that in the form in which we present the proposal here it can be found anywhere in the literature, but we take to be in the spirit of work by Pinkal from the early eighties (Pinkal 1984).

Here the function +, which maps contexts and the contents of formulas in those contexts, onto new contexts, acts as a black box, which fully explicit versions of the theory must spell out in detail.

The main point of the truth conditions in (22) is that they make conditionals like the premises in (13) true. More precisely, let M* = $<C, \leq, c_0, T, +>$ be an SV model enriched with an operation +, and let the set C contain a subset corresponding to the successive stages of a forced march conversation in which the entities $a_1,\ldots, a_N$ play their by now familiar role, i.e. they are introduced one by one as the conversation goes on and, because of tolerance, all end up, one after another, in the positive extension of P, until, that is, tolerance has to bow to some overriding concern. In such a model each of the conditionals $P(a_i) \rightarrow P(a_{i+1})$ will be true in each of the contexts $c_1, c_2,\ldots$ in which the tolerance principle hasn't yet been overruled.

A dynamic setting, with a revised semantics for the conditional as in (22) (and perhaps also for some of the other logical operators) gives rise to its own questions about logical validity and inference. For example, should Modus Ponens (MP) be acceptable given a semantics which includes (22)? An additional complication is that the application of an inference rule may now change the context, e.g. from c to $c+[\![P(a_i)]\!]$. So repeated applications of a rule may lead from sound to overcommitted and, therefore, incoherent contexts (in the extreme case a context in which a_N has ended up in both the positive and the negative extension of P). This is one point where the mentioned accounts that accept (22) part company. In (Shapiro 2005), where (22) is part of a model-theoretic set-up based on Kripke's semantics for Intuitionistic Logic, overcommitted contexts are formally excluded, which is one reason why Intuitionistic Logic, which warrants MP, is a defensible option here. Kamp (1981) does not exclude overcommitted contexts from its formal models, with the consequence that MP is no longer acceptable.[20]

We conclude this subsection by repeating and stressing its main methodological point. The question whether and how logical operators are to be reinterpreted in the presence of tolerance-governed predicates must be clearly distinguished from the contextualist claim that the dynamic context dependence of such predicates is responsible for the paradoxical aspects of forced march conversations. We need an account of what happens in such situations, and why, irrespective of what we then go on to say about Sorites arguments.

2.2.4.2 Non-Dynamic Contextualist Accounts:

Not all contextualist approaches to the sorites are dynamic. There is another group of proposals that is based on a combination of the following two linguistic phenomena. The first is most often observed in connection with nouns, though it seems applicable to predicates of natural language more generally: a noun typically contributes to the noun phrase of which it is the head not just its own, context-independent extension, but a restriction of that extension, which is obtained by intersecting it with a set that can be recovered from context (see section 1).

The second phenomenon is that the contextual restriction imposed on a particular use of a tolerance-governed predicate P can render it immune to forced marches because the restricted application domain contains gaps which exceed the 'span' of the tolerance relation $\sim_P$ that governs P: that is, the application domain of the predicate can be divided into two parts P1 and P2 so that (i) the positive extension of P is included in P1, (ii) the negative extension of P is

[20] Kamp (1981) embarks on an extremely radical revision of the logical operators. Second thoughts on this approach are expressed in Kamp (2013).

included in P2 and (iii) no member of P1 stands in the relation ~p to any member of P2. Let us call such gaps *tolerance-trumping gaps* (*for* a given predicate P with tolerance relation ~p).

The central goal of the proposals that exploit gaps is more modest than that of the contextualist responses to the sorites discussed above. Rather than trying to make sense of uses of vague tolerance-governed predicates in forced march situations, these proposals are aimed primarily at safeguarding uses in which the dangers connected with forced marches are blocked by the presence of a tolerance-trumping gap (thus countering the global pessimism of Dummett 1975, which seems to declare the very notion of observational and other tolerance-governed predicates as incoherent).[21]

A recent proposal of this sort is that of Pagin (2010a,b). Pagin introduces the notion of a **central gap**. The assumption here is that the context in which a tolerance-governed predicate P is used must select one of the gaps in the actual (contextually restricted) application domain of P as 'central gap' and that the borderline between the (actual, contextually restricted) positive and negative extensions of P runs through this gap. In situations with no tolerance-trumping gap in the contextually determined application domain of P, there is no central gap for the context to select. In such cases the use of P is unwarranted. Such cases are beyond the proposal's reach.[22]

One of the apparent consequences of this proposal is that it declares as incoherent even those uses of a tolerance-governed predicate P, in which the application domain of P does contain one or more Sorites series, but where the use of P remains very local, thereby, avoiding the pitfalls connected with those series. Suppose, for instance, that I produce a clearly red colour chip and you describe it as red. Is your description of the chip to be considered incoherent just because it belongs to an actual Sorites series? For instance, I could have the other chips of such a series in my pocket, but you don't know that, and I never produce them or refer to them, so that never become part of the topic of conversation.

Perhaps an advocate of Pagin's proposal would argue that so long as I do nothing to bring the other chips into play this Sorites series isn't part of *red*'s application domain in the context of our exchange. But then the question arises what happens to the application domain when I do produce the other chips. If this does change the application domain, and with it the context that is supposed to determine it, then what happens in the course of a conversation apparently can change the context. And that gets us back to the dynamic dimension of the context dependence of vague predicates.

2.3 Denying the Validity of Sorites Arguments: Changing the Logic via Multi-valued truth or Multiple Notions of Truth

In this section, we look at a few proposals for dealing with the Sorites paradox by adopting a logic in which Sorites arguments are no longer valid. In each of these the change in logic is the result of a revision of the classical semantics of one or more of the operators of first order logic. We have

[21] Wright (1975) can be read as largely endorsing Dummett's position, but in later work he stresses that such a radically sceptical position flies in the face of the fact that speakers do make use of tolerance-governed predicates successully in expressing information and communicating it, and the task of semantic theory is to account for this (Wright 1987).

[22] Earlier proposals in a similar spirit are Manor (1985) and Gaifman (2002). These differ from Pagin's (2010a,b) proposal in that they do make provisions for uses in which the assumption of a tolerance-trumping gap for P is not satisfied. For Manor, all sentences containing P are false in such situations. Gaifman assumes that in such sitations P behaves as a sharp predicate. These assumptions may have certain technical advantages, but it is not easy to see a plausible lignuistic motivation for them.

already come across a few such proposals in passing. In what follows, we look at some further proposals of this kind. We will consider two ways in which such logics can be characterized, that of **multi-valued semantics**, in which use is made of more than two truth values, and that of introducing several **notions of truth** and combining these in definitions of logical validity. We will confine ourselves to one issue which is crucial to what the logics have to say about the Sorites.[23]

2.3.1 Multiple Truth Values.

Many model theories for alternative vagueness logics work with more than 2 truth values. However, as far as known to us, the systems of multi-valued semantics that have been discussed in connection with vagueness are all conservative with respect to classical 2-valued semantics in that their truth value set contains two elements (which we conveniently label '0' and '1') that play the part of falsity and truth and that, more specifically, have the following property: If all atomic constituents of a complex sentence φ have values from the set {0,1}, then φ also has a value in this set, and this value is the same one that we would obtain if we evaluated φ in the 2-valued model M' that can be obtained from M by eliminating all assignments to truth values other than 0 and 1. We restrict attention to multi-valued logics that are conservative in this sense.

Multi-valued semantics offers flexibility for defining consequence relations at two levels. First, additional truth values create room for new definitions of the logical operators, and, second, the possibility for defining logical consequence as necessary preservation of truth can now be varied by allowing for different sets of 'designated truth values': for any proper non-empty subset T of the truth value space, we can define logical consequence as preservation of having some truth value or other within T.

In the spirit of our remarks about conservativity of multi-valued model theories, it is natural to require that 1 must belong to T and that 0 must not. But even then there exists the possibility for more than one definition of the consequence relation even when the total number of truth values is three. An illustration of this is implicit in the part of Section 2.2 devoted to Subvaluationism. There it was observed that the Subvaluationist semantics can be set up as a 3-valued system. Moreover, this set-up can be used to reformulate the partial semantics that we made use of to present the Supervaluationist approach. And when that is done, then, as far as the semantics is concerned, the only differences that remain between Subvaluationism and Supervaluationism are conceptual. But a difference returns when this semantics is used to define a relation of logical consequence. When logical consequence is defined as preservation of truth in partial models, then the resulting logic is Strong Kleene; if consequence is defined as converse preservation of falsity – B is a logical consequence of Γ iff in each partial model in which B is false at least one of the sentences in Γ is false – then the logic is LP.

[23] Both the multi-valued approach to changing the logic and the one which employs different notions of truth are (like those leading to different logics that crossed our path earlier) semantic, consisting of a definition of the truth conditions of the sentences of our language L and a definition of logical consequence that builds on this definition. In the philosophical literature questions about the logic of vagueness are often discussed in terms of the acceptability of particular inference rules, such as Modus Ponens, Disjunctive Syllogism and so on. A general problem with this approach is that it is often difficult to see one's way through to the effects that the alteration of individual rules has for the logic as a whole, both in the technical sense of determining what is deducible, and in the conceptual sense of determining how the over-all concept of validity is affected by such changes. The model-theoretic method, according to which rules and axioms are acceptable if and only if they fit the semantic definitions of logical validity, normally avoids at least the second of these problems: the conceptual motivation for the definition of logical validity is clear from its semantic definition.

Another way to obtain these same two logics within the 3-valued setting is to fix the definition of logical consequence as preservation of designated values while varying the set T of those values: if T is chosen to be the singleton set {1}, the resulting logic is Strong KLeene; if T is chosen to be {1,(1,0)} (where as before (1,0) is the third value), then the resulting logic is LP.

With 3-valued semantics, these are the only two possible choices for T. With more truth values the range of possible options rapidly increases.

2.3.2 Fuzzy Logic

A very different kind of multi-valued semantics and logic is Fuzzy Logic (Zadeh 1965). The truth value space of Fuzzy Logic is the closed interval [0,1] of real numbers. The basic intuition is that the numbers in this interval represent degrees of truth, with 1 representing the maximum degree of complete, unequivocal truth and 0 representing the minimal degree of total, unequivocal falsity. The sense in which these numbers are to be thought of as degrees is visible from Zadeh's truth value rules for negation, conjunction and disjunction.

(23) a. The complement rule for $\neg$: $[\![\neg p]\!]_t = 1 - [\![p]\!]_t$

b. The minimal-degree rule for $\wedge$: $[\![p \wedge q]\!]_t = \min([\![p]\!]_t, [\![q]\!]_t)$.

c. The maximal-degree rule for $\vee$: $[\![p \vee q]\!]_t = \max([\![p]\!]_t, [\![q]\!]_t)$.

The clauses (23b,c) indicate that fuzzy values are not meant to be probabilities. One way to see this is to assume that p has the value 0.5. Then $\neg$p also has the value 0.5. But then it follows from (23b) that $[\![p \wedge p]\!] = [\![p \wedge \neg p]\!] = 0.5$ and from (23c) that $[\![p \vee p]\!] = [\![p \vee \neg p]\!] = 0.5$. In probability logic, these results would be wrong: $[\![p \wedge p]\!]$ and $[\![p \vee p]\!]$ should be 0.5 if $[\![p]\!] = 0.5$, but $[\![p \wedge \neg p]\!]$ should be 0 and $[\![p \vee \neg p]\!]$ should be 1. In fact, these results point towards a deeper and more general problem. There is a strong intuition that $p \wedge p$ and $p \vee p$ should get the same truth value as p – as far as that is concerned, the predications made by (23b,c) are right. But when the arguments of $\wedge$ and $\vee$ are distinct from each other, then, one can't help feeling, the degree of truth of their conjunction or disjunction depends on more than just the truth degrees of the conjuncts or disjuncts. For instance, if the conjuncts are mutually exclusive (in the way that for instance p and $\neg$p are), then their conjunction should get the value of absolute falsity, i.e. 0. No definition for conjunction, as a function from the truth degrees of p and q to the truth degree of $p \wedge q$ can meet all these requirements, because the truth degrees of p and q have nothing to say about any logical connections between them. The same problem arises for disjunction.

Fuzzy Logic has made a virtue of this impossibility of defining binary connectives as degree functions in such a way that all the mentioned desiderata are met by simply setting these qualms aside and introducing various degree-based connectives for special purposes, often in the context of particular applications. In a sense this is also true of the application of Fuzzy Logic to the Sorites Paradox (Goguen 1969). One analysis of arguments like (13) and (12) makes use of the following fuzzy definition of the conditional connective $\rightarrow$:

(24) $[\![A \rightarrow B]\!] = \min(1, 1 - ([\![A]\!] - [\![B]\!]))$

Suppose that $a_1,\ldots, a_N$ form a Sorites series and that $[\![.]\!]$ is a Fuzzy valuation of the atomic sentences of L, which is a monotonically descending function on $\{a_1,\ldots, a_N\}$ in the sense that if i

< j, then $[\![P(a_i)]\!] \leq [\![P(a_j)]\!]$, but in which the difference between successive elements in the series is small. That is, there is some small number $\varepsilon > 0$ such that for all i = 1,..,N-1, $[\![P(a_i)]\!] - [\![P(a_{i+1})]\!] \leq \varepsilon$. Then, with (24) as clause for $\rightarrow$, for all i = 1,..,N-1, $[\![(P(a_i) \rightarrow P(a_{i+1}))]\!] = \min(1, 1 - ([\![P(a_i)]\!] - [\![P(a_{i+1})]\!])) \geq 1 - \varepsilon$. So according to such a valuation, the conditional premises in (2) will all have a high degree of truth – not that of absolute truth perhaps, but certainly a value close to that. Since the first premise is supposed to be unequivocally true and the conclusion unequivocally false, we may assume that $[\![(P(a_1)]\!] = 1$ and $[\![(P(a_N)]\!] = 0$. This means that in this interpretation all the premises have a value close to truth, whereas the conclusion is plain false. That may not be quite as dramatic as when the conclusion is unequivocally false while all the premises unequivocally true. But it still is a reason why (2) should not be logically valid. So what does Fuzzy Logic have to say to that? Again the matter turns on the status of Modus Ponens.

In Fuzzy Logic, there is an aspect to MP and other inference rules that does not come up in other logical systems. Here the 'validity' of an inference rule can be a matter of degree too, in the following sense. The degree of validity of an inference rule $A_1,\ldots, A_n \vdash B$ can be assessed as follows: given any 'minimum' values $M(A_i)$ for the premises $A_1,\ldots, A_n$, what is the minimum of the value for the conclusion that is compatible with any assignment of values $V(A_i)$ to the premises greater or equal to those values (i.e. such that for all I, $V(A_i) \geq M(A_i)$)? For instance, given that the premises A and A $\rightarrow$ B have values of at least a and c, respectively, what is the smallest value b for B such that b and any degree a'$\geq$ a could have yielded a value c' $\geq$ c for A $\rightarrow$ B? Given (24), the calculations are simple for this example. Suppose that A has a degree $\geq 1 - \varepsilon$ and A $\rightarrow$ B a degree $\geq 1 - \delta$; what is the smallest degree b for B that is compatible with these constraints? The answer: $b = 1 - (\varepsilon + \delta)$.[24] Thus, if the premises A and A $\rightarrow$ B of an application of MP have degrees no less than $1 - \varepsilon$ and $1 - \delta$, then we know that the value of B cannot be less than $1 - (\varepsilon + \delta)$. Note that this can be seen as an extension of the sense in which MP is valid in 2-valued logic: In case the degrees of A and A $\rightarrow$ B are 1, and thus $\varepsilon = \delta = 0$, the only compatible degree for B is 1. In other words, the unequivocal truth of A and A $\rightarrow$ B guarantees the unequivocal truth of B.

However, when the degrees of A and A $\rightarrow$ B are less than 1, then, as we have just seen, the degree of B can be less than both of those degrees. Because of this, repeated applications of MP to premises with degrees less than 1 may reduce the trustworthiness of a deduction step by step, so that in the end no constraints on the degree of truth of the conclusion are left. And that is what happens when MP is applied to the premises of (13) as many times as is needed to arrive at the conclusion $(P(a_N)$. The same considerations apply to arguments of the form (12).

Fuzzy Logic thus adds a distinctive new element to the explanation of the apparent validity and actual invalidity of Sorites arguments: Arguments proceeding from less than perfectly true premises may gradually deteriorate, so that eventually they will provide no support for the conclusions reached. (And note well that the justification for classical logic as the correct logic for reasoning with vague predicates, which we mentioned in our presentation of the Supervaluation method, can be retained as well: Classical axioms and rules are verified by the truth definition of Fuzzy Logic, so long as we restrict attention to the extreme degrees 0 and 1. So derivations that are carried out with the help of those rules and axioms can still be regarded as valid in the hypothetical sense that *if* all the premises were perfectly true, then the conclusion would be too.)

---

[24] Let a be the degree of A and c the degree of A $\rightarrow$ B. So (i) $a \geq 1 - \varepsilon$ and (ii) $c = 1 - (a - b) \geq 1 - \delta$. So, substituting $1 - \varepsilon$ for a in (ii) we get: $1 - ((1 - \varepsilon) - b) \geq 1 - \delta$. This can be rewritten as: $b \geq 1 - (\varepsilon + \delta)$.

There are two aspects to the general set-up of Fuzzy Logic and its application to vagueness, however, that can be seen as problematic. The first is the problem we already mentioned about defining the semantics of connectives as functions from and to degrees of truth. That problem arises for conditionals as well: Suppose that the degree of A is 0.5. What should then be the value of ' if A then A' and of ' if A then not A'? What justification is there for insisting that these values must be the same; and how plausible is a clause like (24) in the light of these considerations? The second problem concerns the actual numerical values of the degrees of atomic sentences of the form P(*a*). For these sentences the truth degree clearly derives from the degree to which the argument term satisfies the predicate. But what can be said about the numerical values of the satisfaction degrees of vague predicates?

This of course is an issue that doesn't just arise for Fuzzy Logic. It arises just as much for a probabilistic semantics, in which truth and satisfaction are replaced by probabilities.[25] In neither case is the impossibility to assess values very precisely necessarily a serious objection. Exact degree values don't matter when it comes to articulating and justifying the general principles of Fuzzy Logic. Here certain assumptions about relations between the truth degrees of different sentences will suffice. And even in discussions of particular valuations that reflect our intuitions about Sorites scenarios rough estimates of degrees are often all that is needed, as we did above where we used the parameters ε and δ for 'small numbers > 0'.

Fuzzy Logic was the first framework to offer a degree-based account of sorites problems, but subsequently there have been a number of other degree-based proposals that start from somewhat different conceptual assumptions. Prominent have been approaches that assume a close connection between the satisfaction degrees of vague predicates and degrees of belief, or 'credences' (Edgington 1995, 1997; Smith 2008). These approaches would deserve a discussion in their own right, but for reasons of space we refrain from doing so here.

2.3.3 Multiple Notions of truth

A comparatively recent approach to the semantics and logic of vagueness is that of Cobreros at al. (2010) and Cobreros at al. (2012). Cobreros and his colleagues define three notions of truth for sentences containing tolerance-governed predicates. They do this on the basis of a largely traditional model theory, in which models assign complete extensions to a tolerance-governed predicate P, but contain in addition (sharp) extensions for the tolerance relation $\sim_P$; Cobreros et al refer to such models as 'T-models'. In a T-model M (i) P(*a*) is *classically true* iff $a \in [\![P]\!]_M$ (the extension of P in M), (ii) P(*a*) is *tolerantly true* iff there is some *b* such that $b \sim_P a$ and $b \in [\![P]\!]_M$ and (iii) P(*a*) is *strictly true* iff for all $b \in [\![P]\!]_M$ such that $b \sim_P a$, $b \in [\![P]\!]_M$. These three notions of truth can then be extended to all sentences of the language L by a simultaneous recursion.[26] The

[25] Kamp (1975) offers, as an alternative to Fuzzy Logic in its original formulation, a probabilistic reinterpretation of its truth degrees as probablities. In this alternative the truth clauses for conjunction, disjunction and implication operate on elements of the probability space (i.e. on subsets of the set T of complete extensions of SV models).

[26] Simultaneous in that the strict truth of $\neg\varphi$ is defined in terms of the tolerant truth of φ and the tolerant truth of $\neg\varphi$ in terms of the strict truth of φ. In particular, we get the following clauses for the classical, tolerant and strict truth of $\neg\varphi$:

($\neg$,c) $[\![\neg\varphi]\!]^c_M = 1$ iff $[\![\varphi]\!]^c_M = 0$

($\neg$,t) $[\![\neg\varphi]\!]^t_M = 1$ iff $[\![\varphi]\!]^s_M = 0$

($\neg$,s) $[\![\neg\varphi]\!]^s_M = 1$ iff $[\![\varphi]\!]^t_M = 0$

authors show that a number of familiar logics, as well as some novel ones, can be obtained by means of the standard definition of logical consequence, repeated in (25),

(25) B is a *logical consequence* of Γ iff for every model M, if all A in Γ are true in M, then so is B,

by (independently) varying the interpretation of 'true' in the antecedent and consequent of the conditional in the definiens. When the two occurrences of 'true' are interpreted in the same way, the logics are (more or less) identical to logics that we have already mentioned. On the classical interpretation of both occurrences of 'true' the result is, obviously, Classical Logic; if both occurrences are interpreted as 'strictly true' we get Strong Kleene; and if both are interpreted as tolerant truth we get the Logic of Paradox LP.[27]

Here are some of the things that can be said about forced march situations and soritical reasoning in this setting. Let M be a T-model which describes a sorites situation involving a tolerance-governed predicate P and a sorites series $a_1,\dots, a_N$ for P. Let us assume that the extension of the tolerance relation $\sim_P$ in M consists merely of the identity relation on the Universe of M together with the pairs $<a_i, a_{i+1}>$ and $<a_{i+1}, a_i>$ for i = 1,...., N-1. In M each of the conditionals $P(a_i) \rightarrow P(a_{i+1})$ will be tolerantly true; but one conditional – the one for which i is such that $a_i \in [\![P]\!]_M$ and $a_{i+1} \notin [\![P]\!]_M$ – will be classically false; and both this conditional and the next one, $P(a_{i+1}) \rightarrow P(a_{i+2})$, will be strictly false.[28] Furthermore, we assume that $P(a_1)$ is strictly (and thus classically and tolerantly) true in M (which entails that $a_2 \in [\![P]\!]_M$) and that $P(a_N)$ is strictly (and thus classically and tolerantly) false (which entails that $a_{N-1} \notin [\![P]\!]_M$).

As far as soritical reasoning is concerned, first consider the logic $|=^{tt}$ obtained by interpreting both occurrences of 'true' in (25) as 'tolerantly true'. For this logic, Modus Ponens is not valid. This has to do with the fact that a conditional A → B is tolerantly true iff either A is strictly false or B is tolerantly true. Again, let M be a model in which the line between the positive and the negative extension of P in M runs between a_i and a_{i+1}. Then $P(a_{i+1})$ will be tolerantly true because of the fact that $a_i \sim_P a_{i+1}$ holds in M while a_i belongs to the extension of P. But $P(a_{i+1})$ is strictly false because $a_{i+1} \sim_P a_{i+1}$ and a_{i+1} does not belong to the extension of P; and because

---

For the 'positive' operators &, ∨, ∃, and ∀ tolerant truth of the compound is defined in terms of the tolerant truth of its immediate constituents and strict truth in terms of their strict truth. The truth clause for the conditonal $\varphi \rightarrow \psi$, with its negative occurrence of φ and positive occurrence of ψ, is the one that we obtain by identifying $\varphi \rightarrow \psi$ with $\neg\varphi \vee \psi$; that is, $\varphi \rightarrow \psi$ is tolerantly true iff either φ is strictly false or ψ is tolerantly true; and $\varphi \rightarrow \psi$ is strictly true iff φ is tolerantly false or ψ is strictly true

[27] The qualifier 'more or less' has been put in to distinguish between (i) a first order language L which includes the identity symbol = as well as a relation symbol to denote the tolerance relation and (ii) a language without these two additional relation symbols. The statement in the text is correct as given only for the second, less expressive language. For the first, richer language the options for defining consequence considered by Cobreros et al lead to a finer differentiation of logics. For details see the cited papers.

[28] To see that $P(a_i) \rightarrow P(a_{i+1})$ is not strictly true, consult the truth conditions for $\varphi \rightarrow \psi$ in fn 32 and observe (a) that since $P(a_i)$ is classically true it cannot be tolerantly false and (b) that since $P(a_{i+1})$ is classically false it cannot be strictly true. $P(a_{i+1}) \rightarrow P(a_{i+2})$ is not strictly true since $P(a_{i+1})$ is not tolerantly false and $P(a_{i+2})$ is classically false so not strictly true. But $P(a_i) \rightarrow P(a_{i+1})$ and $P(a_{i+1}) \rightarrow P(a_{i+2})$ are both tolerantly true; $P(a_i) \rightarrow P(a_{i+1})$ because $P(a_{i+1})$ is tolerantly true and $P(a_{i+1}) \rightarrow P(a_{i+2})$ for the same reason: since $P(a_i) \rightarrow P(a_{i+1})$ is tolerantly true, it is not strictly false.

of this $P(a_{i+1}) \rightarrow P(a_{i+2})$ is also tolerantly true. But $P(a_{i+2})$ is not tolerantly true, since none of the entities that are ~P–related to $a_{i+2} - a_{i+1}$, a_{i+2} and a_{i+3} – belongs to the extension of P. This refutes MP for $|=^{tt}$ and by implication for the weaker logics $|=^{tc}$ and $|=^{ts}$, in which the second occurrence of 'true' in (25) is interpreted as 'classically true' and 'strictly true', respectively.

For the logic $|=^{ss}$, on the other hand MP is valid, and by implication the same is true for the stronger logics $|=^{sc}$ and $|=^{st}$.[29] So the argument in (13) can be verified by a valid deduction in those logics. But for these logics the response to the Sorites Paradox is, once again, that in models describing Sorites situations not all the premises are strictly true.

These observations are meant to give a first impression of the range of options that the multiple-notions-of-truth approach makes available. Remarkable about this approach when compared with the other approaches reviewed in this section is the diversity of responses to the Sorites that it covers. Depending on what option is chosen, the resolution of the Sorites either involves the verdict that some premise or premises of the Sorites arguments cannot be jointly true or else that the logic that comes with the given option doesn't support the inference principles needed to derive the conclusions of Sorites arguments from their premises.

That the multiple-notions-of-truth approach offers such a variety of possible responses to the Sorites could be seen as an advantage that it has over the approaches discussed earlier. But there is arguably also a downside to this. More than any of the other approaches we have considered this one leaves us with an *embarras du choix*: One cannot escape the question which of all the different options with which the approach has on offer is the right one.[30]

3. Vagueness and grammar: Additional challenges

Any theory of vagueness will ultimately have to answer the following questions:

- What are the connections between vagueness and morphological gradability?
- What are the connections between vagueness and imprecision?
- What are the consequences of vagueness for the architecture of grammar?

Let us discuss these issues one by one.

3.1 Gradability

There is a strong and widespread intuition that vagueness is connected with gradability. We encountered this idea in the discussion of Fuzzy Logic in Section 2.3.2: What makes a predicate vague is that its satisfaction conditions aren't black and white, but allow for shades of grey; an entity ***a*** can satisfy a vague predicate P to various extents, or various *degrees*, with perfect satisfaction at one end of the spectrum and complete non-satisfaction at the other. So far we have said little about what justifies such intuitions – we only looked at the implications of assuming graded satisfaction for an account of the Sorites. In fact, for all that has been said so far

[29] Suppose that A and A → B are both strictly true. For A → B to be strictly true it must be that either A is tolerantly false or B is strictly true. But if A is strictly true then A is not tolerantly false. So the strict truth of A → B entails the strict truth of B.

[30] Cobreros et al. (2012-2014) themselves elect one specific consequence relation as the right candidate, namely strict-to-tolerant (st) consequence, on several grounds (including for example validation of tolerance, together with maximum preservation of classical inferences, except for the meta-inferential property of transitivity). They also show that one can get rid of the classical notion of truth if one wishes to, and work directly within a 3-valued framework (2014).

justification for this assumption might vary from one type of vague predicate to the next (to the extent that a justification can be given at all).

3.1.1 Relative Adjectives connected with Measuring Procedures

There is one category of vague predicates for which the notion of graded satisfaction conditions has very strong and unequivocal support, and some of these predicates have been prominent in the preceding sections. These are relative adjectives that have a direct connection with measuring procedures, such as *tall, long, short* (all connected with the procedures for measuring length), *heavy, light* (connected with the procedures for measuring weight), *late* (connected with the procedures for measuring time) or *hot* and *cold* (connected with the procedures for measuring temperature).

Measurement of physical magnitudes is a complex topic, with significantly different stories for the different measurable magnitudes,[31] but for purposes of linguistic analysis many of these details do not matter. So we follow the literature on adjectives in simply assuming that with these magnitudes, and with the adjectives that English and other natural languages use to refer to them, are associated *measure functions* f_M, which assign to each entity that can be assessed for the given magnitude M a certain value. We assume that when an adjective is connected with a given magnitude M (in the sense in which *tall* and *short* are connected with the magnitude of length), then the application domain of the adjective consists of those entities that can be assessed for the connected magnitude M. Thus, the application domains of *tall* and *short* both coincide with the Domain of the function f_{length}.

A standard assumption about physical measurement is that functions f_M assign real numbers modulo some *unit of measurement*. Any entity within the Domain of the magnitude can be chosen as unit of measurement for that magnitude. Choosing it means that this entity, and all those for which the measurement function returns the same result, are assigned the number 1 and all other entities a number corresponding to this assignment of 1.[32] For some magnitudes, mostly those that play an important role in our daily lives, more than one unit is in actual use, in large part for the purpose of keeping numerical values within a range with which most of us are conversant (e.g. so that we do not have to deal with very large numbers of millimeters or very small numbers of kilometers). But for the present discussion this variation is of no consequence. So we assume that one unit of measurement is given as part of the measure function f_M and thus that the values assigned by f_M are positive real numbers.[33]

---

[31] For a general discussion of the mathematical properties of measurement see Krantz et al. (1971). For further details about measurement of particlar physical magnitudes one has to consult the branches of science that are devoted to those magnitues.

[32] The correspondences are fixed because each of the physical magnitudes in question has the property that the measure procedures for it reflect proportional relations between the values assigned to different entities (in the sense in which it makes sense to say, for example, that one rod *a* is half the length of some other rod *b* irrespective of what actual numerical length values are assigned to them. Such scales with fixed proportions are known as ratio scales (Krantz et al. 1971). With ratio scales switching from one unit *u* to another unit *u'* is of no theoretical consequence; it simply consists in converting the numerical values according to the unit *u* to the new values (those according to the unit *u'*) by mutiplying them with the factor $f_M(u)/f_M(u')$.

[33] In the standard degree aproach (cf., Venneman & Bartsch 1972; von Stechow 1984a,b; Heim 2001; Kennedy 1999, 2007; Fox & Hackle 2006; Landman 2010; Beck 2011; Solt 2014a, and references therein), the set of degrees is not equated with the set of real numbers, but is assumed to be isomorphic to it. It is thought of as dense,

For adjectives A that have a measure function f_A connected with them there is a very simple answer to a question that we have hardly touched upon so far, viz. the semantics of *comparatives*. In many languages, comparatives are expressed with the help of special morphology that is combined with the morphologically simple form of an adjective A. English has the suffix *–er* to form the comparatives of simple, monosyllabic adjectives, e.g. *tall*, *short*, *smart* as well as a few polysyllabic ones, e.g. *heavy*, while the comparatives of all other adjectives are formed with the help of the modifier *more* (e.g. *more intelligent*, *more educated* and so on). In other languages the details of comparative morphology differ. But a cross-linguistic generalization with few exceptions is that comparative forms are morphologically more complex than the so-called 'positive' forms. This fact suggests that, semantically as well as morphologically, the comparative is more complex than the positive and that it can be obtained from the positive by application of a certain operator -ER, which English expresses either as *–er* or as *more* (Klein 1991; Bobalijk 2012).[34] How can we account for this?

Here is the simple answer: Assume that the basic meaning of an adjective A connected with a measurement procedure simply is the associated measure function f_A (cf., Venneman and Bartsch 1972 and Kennedy 1999, 2007). Then, the semantics of the comparative form of A (to which for simplicity we will refer as 'A-er' irrespective of the form of A) can be given as in (26).

(26) A-ER(a,b) ('a is A-er than b') iff $f_A(a) > f_A(b)$

(26) is appealing for several reasons: It is simple, it is pretty obviously right, and it conforms to the intuition that the semantics of the comparative of an adjective is derived from its basic meaning. But it leaves us with one question and one problem. The problem is that there is only a handful of adjectives that have a measure function associated with them. The vast majority of adjectives that allow for the formation of comparatives do not, and for all of those the story about their comparatives must therefore be a different one. But before we turn to the problem, we first address the question: What does f_A tell us about those uses of A in which it appears in its simplest form, i.e. in its positive uses?

The following answer is typically given within the degree approach (cf., Kennedy 1999, 2007). We must distinguish between the adjective as lexical item and the use of its positive form, just as we have already distinguished between the lexical item and its comparative use. In the case of the positive, this distinction may be less obvious, since morphologically the positive can't be distinguished from the lexical item as such (any more than, say, the 1st person singular present tense form *walk* can be distinguished from the way we identify the lexical verb *walk*; but this analogy can be seen as a warning against taking sameness of surface morphology at face value; cf., Fox 2003). Once the distinction between the adjective as such and its positive form is accepted, the task one faces is to provide a semantics of the Positive as an operator POS which transforms the basic meaning of an adjective A, f_A, into the meaning POS(A) of its positive form.

But what could the semantics of POS be? Kennedy (2007) argues that the semantics of the positive form sets a threshold (standard) for satisfaction of the adjective to a degree that stands

continuous, linearly ordered and supporting operations of addition, substruction and multiplication (von Stechow 1984a,b).

[34] The traditional morphological analysis of adjectives distinguishes three forms: positive, comparative and superlative. There appears to be general agreement that the superlative can be semantically derived from the comparative. That isn't to say that all questions about superlatives have been answered, once an analysis of the comparatives is in place (cf. Teodorescu 2009 and references therein), but in this section we focus on the latter.

out in its application domain, meaning that entities in the positive extension differ significantly from entities outside it with regard to the measurement associated with the adjective. Thus, in effect, this approach appeals to two ways in which the satisfaction conditions of adjectival positives depend on context. First, the context c determines a *standard of comparison* st_c, and second, it determines a *significance margin* ε_c. Assuming that the lexical semantics of A is given by f_A, both st_c and ε_c must be real numbers; and the satisfaction condition for POS(A) can then be given as in (27).

(27) POS(A)(*a*) ('*a* is A') iff $f_A(a) > st_{A,c} + \varepsilon_{A,c}$

The motivations for the two context-dependent parameters in (27) are quite different. The main motivation behind st_c is that the positive uses of relative adjectives are sensitive to the class of objects that is being talked about, within which the adjective is meant to make a significant cut. This intuition seems especially compelling for adjectives whose lexical meaning can be identified with measure functions. The significance margin ε_c is motivated by a different consideration. Consider the following two situations (28a,b) and the descriptions in (29).

(28) a. A picture of two boxes, the left one marginally taller than the right one.
b. A picture of two boxes, the left one substantially taller than the right one.

(29) a. The box to the left is tall, but the one to the right is not
b. The box to the left is tall {as, when, ∅} compared to the one on the right
c. The box to the left is taller than the one on the right

Speakers' judgments about the sentences in (29) in relation to the situations displayed in (28a,b) show a clear difference. While all three of (29a,b,c) are judged correct in relation to (28b), only (29c) is generally accepted in relation to (28a). (29a) and (29b) are felt to be wrong, or false as descriptions of this situation (Kennedy 2007; van Rooij 2011).

Kennedy and van Rooij interpret these results as follows. (i) The comparative statement (29c) is recognised as correct as soon as the slightest excess in height can be observed of the first box in comparison with the second. This is in keeping with the semantics for morphological comparatives given in (26). (ii) The fact that this is not so for the comparison construction in (29b) – that that statement requires a significant difference between the heights of the two boxes, as there is in (28b) but not in (28a) – indicates that the semantics of this construction involves a significance margin that is not required by (29c). Combining this finding with the unacceptability of (29a) in relation to (28a) we can conclude the following. For (29a) to be true, the standard for *tall* must be such that in relation to it the left hand box can be described as (definitely) tall and the right hand box as (definitely) not tall. That requires the standard to lie somewhere in between the respective heights of the two boxes, and so that there is a sufficient significance margin between it and the taller box. When the interpreter is trying to accommodate the standard for *tall* in relation to (28a), she runs into a conflict: on the one hand the standard should be chosen in such a way that the second box does not meet it; that is, the standard must be higher than the height of this box; but since the height difference between the two boxes is so small in (28a), no matter where the interpreter assumes the standard to be located within the interval that separates them, there won't be enough of a difference between it and the taller box.

Precisely *how* the parameters $st_{A,c}$ and $\varepsilon_{A,c}$ are determined by context is left indeterminate by these considerations. Some experimental studies assuming a degree approach suggest that the cutoff st_A in adjective categorization is made at a certain degree or fixed percentage of all the members of the comparison class (Schmidt, Goodman, Barner & Tenenbaum, 2009; Solt & Gotzner, 2012), at the mean or at the mean plus a few standard deviations (Solt, 2011). Additional factors that have been claimed to play a role include the linguistic context of occurrence of an adjective, e.g., the frequency of its co-occurrence with modifiers such as *completely, almost* or *very* (Rotstein and Winter 2004; Kennedy and McNally 2005; Syrett 2007), extralinguistic conventions based on the distribution of degrees among 'normal' comparable entities or among 'normal' stages or parts of an entity (McNally 2011; Toledo and Sassoon 2011), Bayesian reasoning (Lassitier 2011, 2013; Lassiter and Goodman 2013), contextual purposes (McNally 2011; Solt 2012), and subjective evaluations, e.g., by comparison to the evaluator's own degrees (Rips & Turnbull 1980).

The determination of $\varepsilon_{A,c}$ has been argued to depend on salient contextual granularity levels (van Rooij 2011), but has also been claimed to be affected by interests and desires of discourse participants (Graff Fara 2000; Kennedy 2007). The latter are subjective in nature, such that the very same degree difference can be seen as significant in some contexts, but not others. A shift appears to occur from, e.g., quantitative degrees of height to qualitative degrees of desirability of different heights or their fitness to the contextual purpose. We will return to this in the next section with comments on additional factors affecting the determination of the standard and significance margin.

Two concluding remarks to the present section. The first is about gradability and degrees. That adjectives with associated measure functions provide a gradable notion of satisfaction is obvious and trivial: an object *a* satisfies such an adjective A to the degree $f_A(a)$. And the notion of a degree is equally obvious and trivial: The possible satisfaction degrees for A are simply the possible values of the function f_A. We can generalize these notions to the extent it is possible to generalise the notion of 'measure function'. But how can that notion be generalized to adjectives for which no obvious measurement procedures exist, or for, e.g., qualitative scales of, e.g., desirability of heights?

The second remark is about vagueness. According to (26) the meanings of the comparatives of adjectives with measure functions are sharp. Intuitively that seems plausible enough, even though there may be some vagueness in the relation 'taller than'; there are limits to the precision with which the height of any object can be measured, so there are limits to the precision with which it is possible to determine which of two objects is the taller one (or whether, per improbabile, they are exactly equally tall). In this regard, the assumption that the semantics of *tall* is given by the function f_{tall} (= f_{length}) is an idealisation.

But there is, nevertheless, a big difference between the comparative and the positive form of such adjectives. The predicate 'tall', and likewise compound predicates like *tall man, tall basketball player*, *tall six-year old* or *tall building*, are generally perceived as vague, even when they are used in contexts that are fully transparent to the interpreter. How can that intuition be squared with the condition in (27)? The only possibility, it would seem, is to assume that typical contexts, and perhaps all contexts, do not determine the standard and/or significance margin for *tall* completely, but only constrain it within certain limits, even in the case of adjective-noun combinations like *tall man* or *tall building*. For instance, one might, as a first approximation, think of the context c as determining not just two numbers $st_{A,c}$ and $\varepsilon_{A,c}$, but rather two intervals

($st^-_{A,c}$,$st^+_{A,c}$) and ($\varepsilon^-_{A,c}$,$\varepsilon^+_{A,c}$). For an *a* such that $st^-_{A,c} - \varepsilon^-_{A,c} \leq a \leq st^+_{A,c} + \varepsilon^+_{A,c}$, it is then indeterminate whether *a* is an instance of *tall* in this context (for an analysis of the standard as an interval see von Stechow 2009). Again, we are not in a position to say exactly how the context c imposes such partial constraints on $st_{A,c}$ and $\varepsilon_{A,c}$. We will briefly return to this question at the end of the next section.

3.1.2 Relative Adjectives without Measurement Functions

The main problem with the proposal of the last section is that there is only a fairly small set of adjectives that are straightforwardly associated with measure functions (such as f_{length}). Most adjectives, including those with grammatical comparatives, have no such semantic foundation. What can be said about gradability in relation to them?

Attempts to address this question go back to the early history of the theory of vagueness. Kamp (1975) noted that a supervaluation treatment of vagueness creates room for the definition of degrees of satisfaction as ‘probability measures’ over the space of all admissible precisifications of a given ground model. But the paper described this merely as a formal possibility; hardly anything was said about which probability measures among the uncountably many that are formally possible captures the actual semantics of particular predicates (for a promising recent development in this direction, however, see Douven et al. 2013 and Decock & Douven 2014).

Another point in Kamp (1975) is that when an adjective occurs prenominally, then the context in which it is to be interpreted will be determined, wholly or largely, by the noun that follows it. This suggestion was part of the endeavour of that paper to analyse prenominal occurrences of adjectives as vague predicates whose interpretation shows a strong dependence on the context provided by the following noun, rather than as functors operating on the nouns they precede, as proposed in Montague (1968) and other publications from that time. This idea can also be used as the starting point for a definition of satisfaction degrees, but here the story is a good deal more indirect. What follows is an outline of how this story goes.

Klein (1980) makes an explicit proposal for how the nouns preceded by prenominal occurrences of adjectives determine their extensions. The central idea of the proposal is that of a *Comparison Class*. Relative adjectives like *tall*, *heavy*, *impressive* or *trustworthy*, Klein assumes, are typically interpreted in the context of some class X, consisting of entities that are conceived as possible candidates for the application of the adjective. The point of using such an adjective A in the context provided by X is typically to make a meaningful distinction within X, in the sense that both the set of entities in X that satisfy A and the set of entities in X that do not satisfy A are substantial portions of X. Klein captures this intuition by putting the central idea of Supervaluationism, namely that a vague predicate partitions its application domain into positive extension, negative extension and truth value gap, to a new use. In this use, it is the given comparison class X that plays the part that was ascribed in Section 2 to the application domain of A, and the semantics of A is given by the way in which A partitions these different ‘application domains’. That is, it takes the form of a function g_A that assigns to each non-empty set X of entities of which the adjective is applicable in principle, a partition into 1, 2 or 3 subsets.[35]

[35] This is a simplification of the analysis that Klein actually presents. In Klein (1980), the comparison class X is only one aspect of the context. The context may carry further information, e.g. about which of a number of alternative criteria for A is to be taken into account to determine whether A is true of an entity. An example we already

In other words, g_A(X) is of one of three following forms: (i) X or (ii) $\langle X^+,X^- \rangle$ or (iii) $\langle X^+, X^?,X^- \rangle$ (where the sets X, X^+, X^-, $X^?$ are non-empty). The function g_A from comparison classes X to extensions provides information about a reflexive partial order relation $\leq_A$ on the set X: in case g_A(X) = X, all elements of X stand in the relation $\geq_A$ to each other in both directions (so $\geq_A$ is the universal relation on X); in the other two cases – $g_A(X) = \langle X^+,X^- \rangle$ and $g_A(X) = \langle X^+,X^?,X^- \rangle$ – the information about $\geq_A$ is that for each $x \in X^+$ and $x' \in X^-$, $x \geq_A x'$ and not $x' \geq_A x$.[36] As made fully explicit in Van Benthem (1982), natural assumptions about how the values g_A(X) are related to each other for different classes X guarantee that the relation $>_A$ (defined by: $x >_A x'$ iff $x \geq_A x'$ & $\neg\, x' \geq_A x$) is a strict weak order. This means that for any comparison class X the relation $\approx_{A,X}$ as defined in (30) is an equivalence relation and a congruence relation with respect to $\geq_A$.[37]

(30) for all $x,x' \in X$, $x \approx_{A,X} x'$ iff $x \geq_A x'$ & $x' \geq_A x$

Thus, if the relation $>_A$ is used to give the semantics of A-ER, the comparative of A, as in (31) then, modulo Van Benthem's assumptions, A-ER will denote a strict weak order as well.

(31) for all a,b in the application domain X_A of A, A-ER(a,b) iff $a >_A b$

We can make this characterization look even more like (26) by introducing an abstract notion of 'degree of satisfaction of A'. That is, we define the *satisfaction degrees for* A *on* X as the equivalence classes, [x], generated on X by the relation $\approx_{A,X}$ (such that $[x]_{\approx_{A,X}} = \{x' \in X: x \approx_{A,X} x'\}$). Since $\approx_{A,X}$ is a congruence relation with respect to $>_A$, the relation $>'_A$ defined in (32) is properly defined, and it is a strict linear order on the set $\{[x]_{\approx_{A,X}} : x \in X\}$ of satisfaction degrees for A on X (cf., Cresswell 1976). Thus, we can restate (31) as (33), taking the comparison class X to be the application domain X_A of A, and writing '$Deg_A(a)$' for '$[a]_{\approx_{A,X_A}}$' in case $a \in X_A$.

(32) for all $x,x' \in X$, $[x]_{\approx A,X} >'_A [x']_{\approx A,X}$ iff $x >_A x'$.
(33) for all *a,b* in the application domain X_A of A: A-ER(*a,b*) iff $Deg_A(a) >'_A Deg_A(b)$.

The only difference that remains between (33) and (26) is that in (26) the degrees $f_A(.)$ are real numbers, while the degrees $Deg_A(.)$ in (33) are equivalence classes (sets of entities which are

---

mentioned is the adjective *clever*, which can be used as referring to different kinds of abilities, from mathematical problem solving to social adroitness. For more on *clever* see section 3.1.4.

[36] The cases in which the 'division' leads to a single set X do not play any direct role in Klein's proposal. Nonetheless, to capture several cases, the formulation has to include them. One type of case is that where X is a singleton. Obviously it isn't possible to divide a single set into two or more parts. Another type of case involves a larger comparison class X where, nontheless, all elements have to be cassified by A in the same way. For instance suppose that all are of exactly the same height. Then the adjective *tall* cannot make any distinctions within X. Klein's iterative division procedure grinds to a halt when it runs into such comparison classes. Clearly it wouldn't be right to exclude such combinations of adjective and comparison class a priori (cf., Landman 2005).

[37] The equivalence relation R is a *congruence relation with respect to* the binary relation S iff for all x,y such that xRy, if ySz then xSz and if zSy then zSx.

equally A). We could make this look like a perfect similarity by restating (26) also in terms of equivalence classes. In this case, the equivalence classes are those generated by the relation $\equiv_A$, defined by: $a \equiv_A b$ iff $f_A(a) = f_A(b)$ and the linear order >'' between them is given by: $[a]_{\equiv_A} >''_A [b]_{\equiv_A}$ iff $f_A(a) > f_A(b)$ (cf., Bale 2011). But of course that is little more than cosmetics.

There are three morals we want to draw from the Klein-van Benthem proposal. The first is that a partial semantics for vague predicates in the spirit of Supervaluationism can be used as the basis for the semantics of the comparatives of adjectival predicates. All that is needed is to exploit the idea that adjectives are vague predicates which are context-dependent in the sense that their ability to partition their application domain varies as a function of the application domain, as that domain is determined in any given context through the restrictions that the context imposes on it. In this sense, the ordering relations used to interpret comparatives can be seen as emerging from an underlying capacity on the part of the adjective to partition different application domains in different but structurally connected ways. And note that this capacity is the very same that Supervaluationism appeals to in its account of the vagueness of the positive use of adjectives. The difference between Supervaluationism and Klein-van Benthem is that Supervaluationism doesn't exploit the *connections between* the ways in which such vague predicates partition *different* possible application domains.

The second moral concerns the notion of degree. It is a trivial and long familiar fact from the theory of relations that ordering relations with the right properties allow the definition of equivalence relations that are also congruence relations with respect to those ordering relations and that these equivalence relations can then be used to define 'degrees' as the equivalence classes they generate. This enables us to see adjectives supported by measure functions – the degree approach's point of departure – as forming a special case of the more general and abstract construction that we have just gone through.

The third moral arises from juxtaposing the first two. When it comes to the question of the status of degrees in the semantics of adjectives it is important to distinguish between the following sub-questions: (i) Do we need degrees in the semantics of adjectives? (ii) To the extent that degrees are needed or wanted, where do they come from? Do we assume them as primitives – something for which we seem to have good evidence only in those cases in which the semantics of the adjective is rooted in a procedure of measurement – or are they to be defined in terms of something else? And (iii) if degrees have to be justified through definition, in terms of *what* are they supposed to be defined? It is in connection with this last question that the juxtaposition of the first two morals is relevant. A distinction that we believe the literature on vagueness and gradability doesn't always make as clearly as it should is that between definitions of degrees and definitions of ordering relations. Once we have an ordering relation with the right properties, defining a corresponding notion of degree is, we noted, straightforward. But obtaining an ordering relation out of a plausible characterization of the semantics of certain types of predicates, which doesn't build the order into the definition of the predicate in some obvious way from the start, is a less trivial challenge. But it is one that we think Klein and van Benthem have met in an interesting way.

Two questions remain when the approaches of this and the previous section are laid side by side. The first has to do with the interpretation of the positive forms of adjectives. We observed at the end of the last section that positives are vague and that this remains true also when the adjective is followed by a noun that provides the intended comparison class. The account of this section provides a marginal improvement on this in imposing the following requirement: in a context *c* in which the comparison class is the set *X*, the standard $st_{A,c}$ must be such that a significant part of X comes out as satisfying the positive of the adjective and some other substantial part as definitely not satisfying it. But at the same time, we are now also faced with a new problem, which has to do with the significance margin $\varepsilon_{A,c}$.

In (27) $st_{A,c}$ and $\varepsilon_{A,c}$ were both real numbers, and ε_A could be thought of as a significance *interval*. That was possible because for the adjectives discussed in section 3.1.1 degrees were identified with real numbers. According to the generalized notion of degree proposed in this section, however, degrees are in general *not* real numbers, but members of some more abstract order, which, as we have defined it, does not come with an intrinsic metric. Of course, it will still be possible to identify $\varepsilon_{A,c}$ with some degree $>_A st_{A,c}$. But in this present more general setting it is hard to see the difference between this proposal and one according to which the context determines an interval $(st^-_{A,c}, st^+_{A,c})$ such that the positive extension in *c* consists of those members *a* of *X* such that $Deg_A(a) >_A st^+_{A,c}$ and the negative extension of those *a* such that $Deg_A(a) <_A st^-_{A,c}$, whereas the truth value gap consists of the remaining members of *X*.

The second question is related to the first, but it has a more specifically cognitive dimension. Is it right to attribute to adjectives supported by measure functions the role of paradigms for our understanding of adjectival comparatives in general? Do we, as language users, 'model' our semantic conception of adjectives that are not supported by measurement by analogy with those that have such support – that is, by making the tacit assumption that even where a measurement procedure appears to be missing, the semantics is degree-based nonetheless, as if there were some virtual measuring function, that we do not, and perhaps cannot, know? To make this question precise, detailed assumptions are needed about how concepts are cognitively represented and learned. A study relevant to this is Solt and Godzner (2012). Solt and Godzner suggest that changing the degrees of entities matters (affects the standard), even when the ordering between them is preserved. This kind of study can be used to examine adjectives which are not associated with measurement procedures in order to test the generality of the hypothesis that degrees play a fundamental role in their semantics.

Notice however, that a 'degree-less' hypothesis as a response to both questions would still stand a chance to explain the facts as well. Recent developments of the comparison class approach use slightly different constraints to generate ordering relations out of the variance in the extensions of adjectives across comparison classes. These new cross contextual constraints generate semi orders, namely relations such as those denoted by expressions like *significantly older*, *visibly shorter*, or *perceptibly sweeter*. Semi-orders can be used to model the margin of error tolerance parameter (van Rooij 2011b)**.** They are associated with indifference relations (e.g., the relations denoted by *not significantly older*, *not visibly shorter*, or *not perceptibly sweeter*.) Height differences smaller than the perception threshold are rarely relevant for the purposes of speakers in ordinary contexts, rendering such coarse grained relations useful.[38]

The intransitivity of these relations is reflected in their use for *implicit comparison* as in "x is P compared to y", mentioned in the previous section. Such a comparison is only applicable when there is a significantly visible difference between the compared entities (Kennedy 2007). For example, entities x_1, x_2 and x_3 such that x_1 is shorter than x_2 but not visibly so, and x_2 is shorter than x_3 but not visibly so, cannot be described by saying that *x_1 is short compared to x_2* and *x_2 is*

[38] A semi-order > is a binary relation which is irreflexive – $\forall x, \neg[x > x]$, satisfies the interval order constraint – $\forall x_1, x_2, x_3, x_4: [x_1 > x_2 \,\&\, x_3 > x_4] \rightarrow [x_1 > x_4 \vee x_3 > x_2]$, and is semi-transitive – $\forall x_1, x_2, x_3, x_4: [x_1 > x_2 \,\&\, x_2 > x_3] \rightarrow [x_1 > x_4 \vee x_4 > x_3]$. The measurement theoretic equivalence of $x_1 > x_2$ is a function f such that $f(x_1) > f(x_2) + r$, for some constant r representing the perception threshold or margin of error.

short compared to x_3. At the same time, the difference between x_1 and x_3 might already be visible, making it possible to say that *x_1 is short compared to x_3*. Similarly, if an equative relation like *as short as* can be used to express indifference, the statements *x is as short as y* and *y is as short as z* can be rendered true when the statement *x is as short as z* is false (van Rooij 2011b). Thus, according to van Rooij (2011b), explicit comparatives such as, e.g., *John is taller than Mary*, should be analyzed as weak orders (via van Benthem's constraints), while implicit ones, e.g., *Compared to Mary, John is tall*, should be analyzed as semi-orders (via van Rooij's constraints), to capture their insensitivity to small distinctions.

Importantly, the indifference relation $\approx$, which is defined in terms of a semi-order > by: $x_1 \approx x_2$ iff $\neg[x_1 > x_2]$ & $\neg[x_2 > x_1]$, does not divide the domain into equivalence classes (it is reflexive, symmetric but not transitive). Thus, there is no numerical representation f_P such that for every x_1 and x_2, $x_1 \approx x_2$ iff $f_P(x_1) = f_P(x_2)$ (van Rooij 2011b). In other words, when in a given context c, an adjective imposes upon its application domain the structure of a semi-order, its semantics cannot be reconstructed as a degree semantics. A degree semantics is possible only in a different context c' in which differences are recognised that are ignored in c. That is, in c' pairs $\langle x_1, x_2 \rangle$ of elements of the domain which count as indifferent in c are treated in c' as ordered (i.e. $x_1 > x_2$ or $x_2 > x_1$). In this way the semi-order of c can become a weak order in c' and from this weak order it will then be possible to construct degrees in the now familiar way. For example, consider a context in which pairs which differ in height by less than 10 millimeters are loosely considered *equally tall*. If these pairs are added back into the entity set denoted by *significantly taller* (and *less tall*), the result is a strict weak order, and the respective indifference relation is transitive.

In fact, there is a unique largest weak order $>_P$ extending a semi-order >, and it is this weak order that is relevant for derived comparatives: x_1 is $>_P$ x_2 iff $\exists x_3$ $[(x_3 \approx x_1$ & $x_3 > x_2) \vee (x_3 \approx x_2$ & $x_1 > x_3)]$, e.g., x_1 is strictly speaking taller than x_2 iff x_1 is indistinguishable from a x_3 who is taller than x_2, or x_2 is indistinguishable from a x_3 who is shorter than x_1; x_3 is the *witness*. Thus, equivalence relations can be derived systematically to generate degrees when these are needed (e.g., upon the use of numerical measure phrases). However, degrees are assumed to not play any role in constructions within which nothing motivates their derivation, such as positive forms like 'Dan is happy', and gradable constructions involving merely *-er* and *very*.

Van Rooij's observation that some adjectives have, in some contexts at least, a semantic structure that prevents them from being analyzed as degree predicates is important because degrees are often assumed to be indispensible for a compositional account of these and other gradable constructions (for a well developed view of the syntax-semantics interface see Heim's work in, for example, Heim 2001 and as reviewed in Beck 2011; see also Landman 2010, a.o.). However, analogous accounts within the comparison class approach have recently emerged as well (see discussion of the syntax-semantic interface in a degree-less account in Doetjes, Constantinescu and Součková 2010). Instead of using degree variables, these accounts make use of comparison class variables in logical forms (as well as expressions of more complex types involving such variables). Thus, as in the degree approach, a difference is assumed to exist between the meaning of an adjective (a function from comparison classes to extensions) and the meaning of its positive form. The latter effectively provides a comparison class variable for the adjective to operate on in order to yield a set of entities.

If it isn't always possible to analyze adjective occurrences as degree predicates, and yet there are many occurrences whose form and meaning strongly suggest such an analysis, is there a way to reconcile these constraints? One possibility that deserves to be looked into more closely is that of an interpretation strategy involving some form of coercion, along the lines of Partee (1987). In that paper, Partee is concerned with the question whether a semantic theory should always ‘generalise to the worst case’. No, she argues, often a more elegant approach is to take the simplest interpretation to be the basic one, while adopting a highly constrained system of type shift operations to derive more complicated interpretations. The situation before us is different – worse, you might say, because there isn’t even the possibility of generalizing to a worst case. But a similar approach may well be right here as well: adjective occurrences whose morpho-syntactic environment does not require a degree analysis are treated by default without an appeal to degrees; but allow for shifting to a degree-based analysis for those occurrences that imply a need for it. The former treatment would assume no more than a semantic structure with the properties of a semi-order, the latter would require a weak order structure at a minimum and in some cases the degrees that can be constructed from this weak order may coincide with the values assigned by an associated measure function.

Following this line of reasoning, Booij and Sassoon (2013) hypothesized that adjective categorization decisions are made on the basis of ‘degree leaps’ - in such a way that language users use the coarsest semi-order that is compatible with the interpretative task that the occurrence of an adjective confronts them with. Only when this semi-order is too coarse to split the given comparison class into two subsets (X^+ and X^-), will they look for a finer semi-order (and thus for smaller leap functions) to give them a viable cutoff point. The results of an experiment with adult participants supported this hypothesis. However, this was not so when the experiment was done with 8-year-olds. Here the results varied, with little agreement. The differences between the results obtained with adults and those obtained with the children suggest that the situation is more complex than was initially assumed.

To explain this, here are more details of settings such as those used in the experiment. Imagine a set of pencils of 1 cm, 3 cm and then a linearly growing set, 4, 5, until 15 cms. The biggest leap exists between the two smallest items; yet intuitively, the 3 cm pencil is not big. Hence, probably, adults classify objects using a conjunction of rules such as finding a leap and seeing that the leap is in a "sensible" place, roughly above the middle of the scale (possibly because this position leaves room for a leap also at the boundary of the antonym extension). Children appear to use only one of the two rules at a time, plausibly because they still have difficulty with conjunction rules, as well as with access to unmentioned alternatives such as the antonym of the adjective. Thus, some of them placed the cutoff in the lower leap, some placed it in the upper part of the scale and some placed it in the middle of the scale.

What can we conclude from these experimental findings? One result suggested by this and other studies (cf., Tribushinina 2013) is that there is a tendency for adults to interpret positives of adjectives more restrictively than do children. But the propensity toward restrictive interpretations appears to be malleable. For instance, Panzeri and Foppolo (2011) found that after adult subjects had been trained to not derive scalar implicatures (i.e., to accept the use of existential statements in contexts at which the corresponding universal statements are true), they tended to choose lower standards for adjectives. Apparently weakening a human interpreter’s propensity to put maximally strong interpretations on utterances in one domain can have the effect of weakening that propensity also in other domains.

Other acquisition studies show that the positive forms of adjectives are acquired earlier than forms whose interpretation involves degrees, and that, among the latter, those involving mere gradation are acquired before those involving a metric structure (as in *10 cm taller*). This order in acquisition is echoed by typological findings (Ravid et al 2010; Beck et al 2010). Some languages posses no measure phrases, but do possess comparisons or at least positive forms (Beck et al 2010; Solt 2014a).

These developmental and cross-linguistic findings make it tempting to think that the coarse-grained interpretations of positive forms are at the bottom of some kind of hierarchy of increasingly complex interpretation strategies for different morpho-syntactic constructions in which adjectives can appear. But we have to be careful with such speculations. It may well be that the adult speaker of a language like English, who has mastered all these strategies has a different semantics for positive adjectives than the child that is on the bottom rung of this developmental ladder. Possibly it is by climbing the ladder that the language learner arrives at the mature understanding of forms for which he needed some different initial command to set him of on his ascent.

3.1.3 Absolute Adjectives

Adjectives differ in various ways. For one thing, not all adjectives admit of comparatives, or admit of them happily. Examples of those that don't are *spherical*, *valid*, *four-legged*. But also among those adjectives for which the comparative form is fully acceptable significant differences can be observed. The difference on which we focus in this section is that between relative adjectives, like *tall*, *heavy* or *smart*, i.e. those that were the topic of sections 3.1.1 and 3.1.2, and *absolute* adjectives like *flat* or *wet*. (34) reveals some ways in which the latter are judged to behave differently from the former, but these examples also show that *flat* and *wet* differ from relative adjectives in different ways.[39,40] (34b) shows that, conceptually at least, there is a maximum to the degrees to which something can be flat, and that the positive form *flat* indicates that maximal degree. When something qualifies as flat, then there can't be anything else that qualifies as flat to an even higher degree. In that sense *flat* is 'absolute'. On the other hand, the acceptability of (34a) indicates that 'not flat' is not used to describe a maximal degree of non-flatness, presumably because no such maximum is conceptually available. With *wet*, things are the other way around. There is no maximal degree to which a thing can be wet, which is why (34d) is unproblematic. But 'not wet' is absolute: If something is claimed to be not wet, then it is inconsistent to add that something else is even less wet (i.e. that the first thing is wetter than that second thing, see (34c)). The degrees to which things can be wet can be seen as 'bound from below', although in this case the bounding degree is not a degree of wetness, but the maximal degree of being not wet, or dry. As shown by the acceptability of both (34e) and (34f), the degrees of a relative adjective like *tall* are bound in neither direction.

(34) a. The desk is not flat, but it is flatter than the pavement.
b. # The pavement is flat, but the desk is flatter.
c. # The desk is not wet, but it is wetter than the pavement.
d. The pavement is wet, but the desk is wetter.
e. Bill is not tall, but he is taller than Mary.

[39] Speakers vary in their judgments regarding the examples in (34), thus we are not sure how reliable and widely shared these judgements are. What further complicates the data are variations in the formulation of these sentences. For instance, both (34a) and (34d) seem to improve when *flatter*/*wetter* are replaced by *even flatter*/*even wetter*. However, we will leave these problems aside and continue on the assumption that the assessments in (34) are correct.

[40] And recall the examples in (4) of section 1.1, which also show differences between absolute adjectives of the *flat* type and those of the *wet* type.

f. Bill is tall, but Mary is taller.

Adjectives of the type of *flat* are called *total absolute* adjectives and those of the type of *wet* are called *partial absolute* adjectives.

It should be emphasised that the distinction between relative adjectives and the two kinds of absolute adjectives exemplified by *flat* and *wet* is ultimately a conceptual one. The distinction may have something to do with the way we understand the world in which we live. For instance, it may be part of that understanding that there is no upper bound to how tall a person or a building can be, in the way that we do understand that there is a limit to how flat a surface can be, which is reached when it has no unevennesses of any sort. But there clearly is a physical limit to how 'unheavy', i.e. light, a physical object can be – viz. when it has zero weight, and yet that absolute lower bound is not part of the concept 'heavy'. Likewise, it has been widely known for well over a century that there is an absolute zero point to temperature, but the adjectives *hot* and *cold* are relative adjectives that by the tests in (34) behave like *tall* and not like *flat* or *wet*.

In the remainder of this section we focus on total absolute adjectives.

The semantic connection between the positive and the comparative forms of a relative adjective *A* that we discussed in the last section was this: the relation expressed by the comparative *A*-er holds between two objects *a* and *b* in the application domain of *A* iff the *A*-degree of *a* is bigger than the *A*-degree of *b* (however *A*-degrees may be defined or construed); and in a given context c, *a* satisfies the positive POS(*A*) of *A* iff the *A*-degree of *a* equals at least $st_c(A) + \varepsilon_c(A)$ (with $st_c(A)$ the standard of A in c and $\varepsilon_c(A)$ *A*'s significance margin). Since relative adjectives involve neither maximal nor minimal degrees, there will, for any context c, be both bigger *A*-degrees and smaller *A*-degrees than $st_c(A)$, and the same will be true for $st_c(A) + \varepsilon_c(A)$. For a total adjective *A* the relation between the positive and the comparative forms of *A* is clearly different. Now the degree scale of *A* has a maximal element and that, intuitively, is the standard that *a* should live up to in order to qualify as satisfying POS(*A*).

It has been stressed, in particular by Unger (1975), that when this is taken literally, it puts the satisfaction level for POS(*A*) at an impossibly high level, which can never be attained in actual practice. If we were to stick to such levels, the positive forms of absolute adjectives would be useless: any assertion of the form '*a* is *A*' would be false. As it is, we do make such statements: 'This desk is flat', 'This towel is dry/not wet', 'This glass is full/empty' and so forth and we succeed in conveying genuine information by them. What is it that makes this possible?

Intuitively, the answer would seem to be obvious, especially in light of what has been said about the positives of relative adjectives. The context affects the interpretation of absolute adjectives, though the way in which it does that differs from how it affects the interpretation of relative adjectives. Its primary role is to fix the standard of *A* in a 'realistic' manner, which objects in the application range of *A* have a chance of meeting. We can think of this element of context dependence as given by the lower bound $lb_c(A)$ of a tolerance interval whose upper bound is the context-independent standard Max(*A*). In other words, *a* qualifies as an instance of the positive form of *A* in c iff the *A*-degree of a lies within the degree interval $(lb_c(A)$,MAX(A). If *A* is associated with a measure function f_A, then this can be stated as in (35).

(35) (POS for total absolute adjectives A)
$\text{POS}(A)(a)$ iff $lb_c(A) \leq f_A(a) \leq \text{MAX}(A)$

But this cannot be all. If justice is to be done to the observation that (34b) is ill-formed, then the determination of $\mathrm{lb}_c(\mathrm{A})$ also must have a further effect: The ordering imposed by f_A should be made inaccessible to the semantics for the values of f_A in the closed interval $[\mathrm{lb}_c(A),\mathrm{MAX}(A)]$. One way to do this is to count all objects *a* such that $f_A(a)$ belongs to $[\mathrm{lb}_c(A),\mathrm{MAX}(A)]$ as having the same A-degree. This renders sentences like (34b) trivially false (cf., Rotstein and Winter 2004).

These stipulations make the semantics of the positive and comparative forms of total absolutes come out the way we want. But as things stand there is a lack of uniformity between these principles and the ones put forward in the previous sections for the positives and comparatives of relative adjectives. Perhaps the difference between these two adjective types just doesn't admit of a more unified treatment than the one we have sketched.[41] However, the possibility of a unified account should not be given up too easily. We leave this for further reflection, and conclude the discussion with one last comment on the potential of looking at the role of dynamic shifting of the granularity level used within contexts.

One way to think of the 'loose' interpretations of the positive forms of absolute adjectives is in terms of partitions of associated degree scales into segments. For a total adjective, the partition should be such that its positive extension consists of those objects that fall within the highest segment of the partition – the one that has the maximum of the scale as its upper bound. Suppose, for example, that the semantics of the total adjective *flat* is given by a flatness scale. This scale can be partitioned into adjacent segments in various ways. Some of these will be coarse-grained, involving only few largish segments, others more fine-grained, with more but smaller segments. Part of what is involved in the interpretation of an occurrence of *flat* is to choose a partition of the scale or to find out from the context what the intended partition is. Positive occurrences of the adjective will show a preference for coarse-grained partitions. In order to satisfy the positive form, the object of the adjective being predicated must fall within the 'top' segment, and that will only be the case if this segment is big enough. Thus, for instance, for the sentence 'The pavement is flat' to come out as true, the partition must be such that the pavement has a degree of flatness that falls within its top segment. To accommodate what they perceive to be the intentions of the speaker, interpreters will often go for such coarse-grained partitions, thereby justifying the speaker's claim.

But while positive occurrences of a total positive adjective push us in the direction of coarse-grained partitions, occurrences in comparison constructions push us towards fine-grained ones, which can capture the degree distinctions that are involved in the semantics of these constructions. For instance, the statement 'The desk is flatter than the pavement' may be true, and recognized as such, even though the difference between them is only slight.

This tension between the need for coarse-grained partitions for the interpretation of positive forms and for fine-grained partitions for the interpretation of comparison constructions manifests itself dramatically when the two conflicting needs arise in connection with a single sentence like

[41] Such a conclusion would be consistent with the proposals of van Rooij (2011a,b), which deal with total absolute adjectives by means of different *cross-contextual* constraints than those that are assumed for relative adjectives. For total absolute adjectives, van Rooij uses constraints drawn from Arrow (1959), which guarantee that they denote functions from comparison classes to their highest-ranking members under the ordering denoted by their derived comparatives. Burnett (2014) provides different cross-contextual constraints for absolute adjectives, as well. In addition, she rules out orderings within the positive extension of total absolute adjectives by fixing their comparison class as the entire domain of individuals. Variability in the comparison class is only introduced when a total absolute adjective is interpreted tolerantly. It is the context relative imprecise tolerant interpretation of absolute adjectives that is gradable.

(34b). Here the truth of the first conjunct requires a coarse-grained partition which enables the pavement to make into the top segment, whereas the second conjunct requires a partition that is fine-grained enough to draw a distinction between the pavement and the desk, and one that is such that the pavement does *not* end up in its top segment. Clearly a single partition cannot deliver on both fronts. It is our impression that if the two conjuncts were further away from each other, for instance, if they were separated by another clause or sentence, then the awkwardness of (34b) would be mitigated, but more empirical work is needed to confirm this.

A more general question for empirical research is how interpreters navigate between partitions of different granularity when interpreting the different forms of the different types of adjectives (for relevant data points and different theoretical perspectives see Unger 1975; Kennedy 2007; Syrett 2007; Rotstein and Winter 2004; McNally 2011; Sassoon 2012).[42]

3.1.4 Additional types of adjectives

The distinction between relative and absolute adjectives is not the only one that is relevant for a semantic theory of adjectives. Adjectives such as *tall* and *heavy* are *one-dimensional* in that their satisfaction degrees are given by a single measurement function – height for *tall* and weight for *heavy*. But such adjectives are comparatively rare. Far more often, there are several criteria that are relevant to how well an object instantiates a given adjectival predicate (Kamp 1975, Klein 1980, Murphy 2002, Sassoon 2013a,b). A classical example from the literature is *clever*: whether someone is to count as clever, and how clever, may depend on different kinds of skills, from mathematical problem solving to such social skills as persuasiveness or successful manipulation of others (see Grinsell 2012 and Canonica 2014 for discussion of vagueness in sets of criteria used in classification under adjectives).

For an adjective like *clever*, the context is important for more than one reason – not only because it may fix the set of categorization criteria and the standard of application for either or all criteria, but also because it determines the way in which the different criteria integrate to create a single categorization criterion. The dimensional criteria can be conjoint or disjoint; alternatively, context may single out one of them as the only one that is relevant, or it may impose a certain weighting between them. This kind of 'multi-dimensionality' is extremely common. We find it also, for instance, with *flat*: How do we compare the flatness of a very smooth flat surface with one small bump and that of a surface that has a pattern of tiny ribs all over it? Which of the two is the flatter one?

One reason why the difference between one-dimensional and multi-dimensional adjectives is important is that they have different 'logics'. For a one-dimensional adjective, the positive is

[42] We may add at this point that our discussion of absolute adjectives abstracts away from many additional factors that enter into our use of many such adjectives. An example is *full*. McNally (2011) proposes that classification under an absolute adjective is rule-based, but different social conventions dictate what counts as, e.g., *full* in different contexts (McNally 2011). When wine is served, the glass is normally only half filled, when Whiskey is served, it is not even half filled, and even when water or tea are served, the glasses are hardly ever filled to the brim. By contrast, classification as a relative adjective, on this view, depends on no rule at all. Rather, it depends on similarity to other concrete examples (e.g., a pencil slightly shorter than a big pencil is considered big as well). Nonetheless, the rules for classification under absolute adjectives do introduce *social and subjective variability* into interpretations; e.g., in the context of a restaurant, a half filled glass will sometimes count as full, and sometimes not, depending on, e.g., the type of glass and the liquid served. A client who is served a half filled wine-glass cannot ask his or her money back, for such a wine glass *does* count as full (McNally 2011).

dependent on context, but the comparative is not. For such an adjective *A*, it will invariably be the case that if *a* and *b* are both in the application domain of *A*, then

(36) either a is A-er than b or b is A-er than a or, thirdly, a and *b* are equally *A*.

For multi-dimensional adjectives the comparative will in general depend on context as well as the positive. The context may influence the comparative of a multi-dimensional adjective by, for instance, selecting one of its dimensions rather than another. If the context doesn't provide the right information, then the disjunction in (36) need not hold.

What we see here is a further twist to the 'logic of vagueness' discussed in Section 2: It isn't just vagueness as such that raises questions about logic (that is, about the logic of languages with predicates that are vague in whatever way and for whatever reason). The questions we have just raised concern the vagueness-related semantic structure of particular types of predicates and the logical implications of such structures. Some such predicates have a comparatively simple semantic structure, which supports logical laws and inferences that do not hold for predicates with a more complex structure. But those with more complex structures may licence other logical principles in their turn. Not all that much seems to be known at this point about how the semantic structure of vague predicates may vary and what implications that may have for the logic of those predicates. Further distinctions have been noted in the literature between different types of adjectives, but we cannot go into more detail about this here (see Kamp and Partee 1998; Morzycki 2014; McNally, this volume, and references therein). But the topic is important, and it is important in particular for natural language predicates that take the form of adjectives. Here the relations between meaning, logic and form have a special linguistic importance, since in many languages adjectives can occur in a rich variety of morpho-syntactic constructions that reflect their semantic and logical properties (positives, comparatives, equatives and more). The study of the variety of semantic types of adjectives is therefore as important for the linguist as it is for the logician. But this point extends far beyond the category of adjectives. For a discussion of different types of nouns and verbs in a similar range of constructions see, for example, Wellwood, Hacquard and Pancheva (2012) and references within.

We have just argued that questions about the logic of different types of vague adjectives (as well as of other types of vague expressions) differ from those that were touched upon in the discussions of the logic of vagueness in Section 2. In fact, investigations of the logic of different types of adjectives and other vague predicates should *build on* a general logic of vagueness (whatever that logic is taken to be). The logical principles that are specific to particular types of predicates should be cast in such a form that they can be incorporated into such a general logic, for instance as axioms that are restricted to the predicates of the type in question.

3.2 Imprecision and Approximation

Our discussion of total absolute adjectives like *flat*, *dry* or *clean* suggests the following picture: these words are predicates with a well-defined semantics, but one which in practice no object can ever meet (e.g., nothing is ever completely flat). So in order to be able to make a meaningful use of such a predicate *A* we always relax the degree to which they should be satisfied, and this relaxation varies as a function of context: In any such context *c,* the 'de facto extension in *c*' consists of all objects whose satisfaction degree lies within the interval $[lb_C(A),MAX(A)]$.

But absolute adjectives are not the only expressions for which an account of this sort suggests itself. Some others are on display in (37) (Lasersohn (1999), Krifka (2002)).

(37) a. Mary arrived at three o'clock.
b. Mary arrived one minute after three o'clock.
c. The distance from Amsterdam to Vienna is 1000 Km.
d. The distance from Amsterdam to Vienna is 965 Km.
e. In decimal notation π is 3.14.
f. In decimal notation π is 3.1415926538.

The time denoting NPs *three o'clock* in (37a) and *one minute past three o'clock* in (37b) are often, and plausibly, said to denote instants of time, which occupy infinitesimal parts of the timeline. But if that is so, then it is hard to see how (37a) or (37b) could be really true. For how could the events described in those sentences be that narrowly confined in time? So, if there are to be situations that these sentences describe correctly, then that must either be because the relations between the described events stand to the denoted instants in some other relation than temporal identity (i.e. identity of the time of the event with the denoted instant); or else the temporal adverb must make available something other than an instant, such as an interval surrounding it, within which the event *e* must be temporally contained. We do not know how much there is to choose between these options, but we follow Lasersohn (1999) and Krifka (2002, 2007) in pursuing the second.

According to this second option each context *c* for the use of *three o'clock* must determine an interval, and it seems obvious that that interval should include the instant *t* denoted by the phrase. A further natural assumption is that *c* determines an interval that has *t* as its centre point. That is, *c* determines a temporal distance ε_c such that (37a) is acceptable in c only if the described event is temporally included in the interval ($t-\varepsilon_c$, $t+\varepsilon_c$).

(37b) is like (37a) except that we tend to allow for more tolerance in the case of the 'round' time-denoting phrase *three o'clock* than with the phrase *one minute to three o'clock*, which is 'not round'. For instance, when nothing more is said and Mary arrived between two and three minutes past three o'clock, then (37b) would normally be regarded as wrong, but (37a) might well be considered acceptable.

Sentences (37c-f) serve to illustrate much the same points as (37a,b). Krifka (2002) notes that if the distance between Amsterdam and Vienna is in fact 972 KM, then (37c) might be accepted but not (37d), in spite of the fact that the discrepancy between the actual distance and the one mentioned in (37d) is less than that between the actual distance and the one mentioned in (35c). The sentence pair (37e,f) illustrates another instance of the same phenomenon. 3.1415926538 is closer to the true value of π (of which the first 10 decimals are 1415926535 rather than 1415926538) than it is to 3.14. Yet our convention about decimal approximations qualify (35e) as right but (37f) as not: What matters according to that convention is that those decimals must be correct which are explicitly given. In science this principle governing the interpretation of decimal approximations is generally accepted and rigidly observed. The principle that round numbers guide us towards more tolerant contexts is less rigid and not as easily stated, but it is in the same spirit.

In our comments on the examples in (37) we have been anxious to avoid the use of the words 'true' and 'false'. For instance, we didn't say that (37e) is true and (37f) false, though that might be thought to be a perfectly correct description of the difference between them. In avoiding 'true' and ‘false’ we have tried to remain consistent with Lasersohn’s discussion of these phenomena, according to whom, for instance, both (37a) and (37b) are equally false when Mary arrived at two minutes past three, even if the context places the event *e* of Mary’s arrival within the ‘halos’ of both *three o’clock* and *one minute after three o’clock* (where the halo of *three o’clock* in context *c* is the interval around three o’clock within which Mary’s arrival must fall in order for (37a) to be an acceptable description in *c* of the given situation.) But in light of the discussion throughout this paper, we cannot see any compelling reason against describing the sentences in (37) as true or false, so that in a context *c* which includes the event *e* within the halo of *three o’clock* but not within that of *one minute after three o’clock*, (37a) will be true and (37b) false (among representatives of this perspective are Lewis 1979, Rotstein and Winter 2004 and McNally 2011). Consider once more (37a,b). We can treat (37a) as a special case of the semantics of vague predicates that has been discussed throughout this article by assuming that when the prepositional phrase *at three o’clock* is used as an adjunct to an event description in a context *c*, then it is interpreted as an event predicate that is satisfied by an event *e* iff *e* is temporally included in the halo of *three o’clock* in *c*.

Krifka accounts for the ‘round number’ effect as the result of a competition between the demand for precision, on the one hand, and a preference for short expressions on the other: The phrases for round numbers tend to be simpler and shorter than those for numbers that are not round, and the advantage of brevity will often outweigh the striving for precision. It is not clear, however, that brevity is always the reason why a speaker chooses a round number phrase instead of a longer but more accurate one. It isn’t just that round number phrases come with an implication of greater tolerance – something that only involves a ‘halo-ish’ interval around their precise denotation – but, more globally, their use relates to a certain coarse granularity that extends to the entire numerical domain. For instance, the use of the phrase *1000 Km* in (37c) suggests a context in which distances are assessed as multiples of, say, 50 Km intervals (see Krifka 2007 for an elaborate discussion).

Canonica (2013) argues that such coarse granularities are advantageous from a computational perspective, especially in relation to statements that involve quantification over the given numerical dimension. Shifting to a coarser granularity makes it possible to switch from infinite to finite domains when evaluating such statements. In a similar vein, Solt, Cummins and Palmović (2014) argue that coarse and approximate interpretations reduce processing costs. Vagueness and approximation are useful for reasons of politeness and diplomacy as well, e.g., describing a strictly speaking *incorrect* utterance of a speaker as *roughly correct* can help to save his or her face. Fixing an appointment for *the afternoon* is less committal than fixing it for *4:00 o’clock,* meaning that the chance one has of arriving ‘on time’ increases (van Deemter 2010). Rather than being viewed as mere deficiencies of natural language expressions (Frege 1903; Russell 1923), vagueness and imprecision, on these approaches, have a purpose.

Perhaps most importantly, speakers themselves often have only approximate knowledge of the facts they want to communicate. I do not know exactly how far it is from Amsterdam to Vienna, but I do know it is about 1000 Km. I could of course convey this approximate knowledge by saying something like ‘The distance from Amsterdam to Vienna is roughly 1000 Km’. But in view of the round number conventionm (37c) will do just as well, and in the choice between these two potions brevity can easily have its way, for in this case no alternative benefits are competing with it. Furthermore, even when I do know that the distance from Amsterdam to Vienna is 972 Km, I may nevertheless prefer to say ‘The distance from Amsterdam to Vienna is

1000 Km', as a way of conveying that the imprecise information this formulation conveys is all that I take to matter.

To conclude this section, is imprecision of the kind illustrated by the sentences in (37) a form of vagueness or isn't it? Perhaps it is impossible to answer this question without some measure of legislation. It should be clear from what we have said here that we see imprecision of this kind as a form of vagueness. One justification for this that we haven't drawn attention to so far is the following: this kind of imprecision shares some of the features that we find with prototypical examples of vague predicates, such as *bald* or *heap* and that are seen by many as the hallmarks of vagueness: giving rise to sorites problems – if *a* belongs to the halo of *x* and *b* is only marginally different from *a*, how could *b* fail to belong to this halo? There are means of contextual resolution (through contraction of halos), but without the chance that this will ever ban the dangers of the Sorites altogether.

3.3 Vagueness as an aspect of Grammar?

We have discussed a variety of adjective types (cf., section 3.1), as well as other expressions (cf., sections 1.2 and 3.2), that manifest vagueness. One way of summarizing these observations is to say that natural language grammars are affected by vagueness. In fact, that grammar should be 'vague' in this sense is nothing new. Various suggestions have been made, both with reference to (potential) lexical items and to certain syntactic constructions, that it is vague whether the item is part of the language, or whether it can be used with a certain argument structure, or whether the construction is really part of the language or only something that people will say when they don't pay attention and thus are prone to making mistakes. For such reasons grammar has long been seen by many as inherently 'fuzzy' (cf., Aarts 2004 and multiple references within).

In fact, such fuzziness seems to go hand in hand with language change. To give just one anecdotal example, one of us still remembers being somewhat bewildered quite a few years ago when hearing for the first time an American waiter speak the word(s) 'Enjoy!', evidently as an encouragement to enjoy the food he had just placed on the table. One gets used quickly enough to such innovations, but even today *enjoy* doesn't seem a very good paradigm of an intransitive verb. 'After they enjoyed, the men withdrew to the smoking room for port and cigars' seems a strange way of saying 'after they enjoyed the meal', and not just because of a clash of cultures. On the whole, grammaticalisation is a complex process, involving slow and tentative progress through many gradations of rejection, uncertainty and acceptance. It seems plausible to say that something like vagueness or indeterminacy governs the intermediate stages of such processes.

The case for a deep interaction between vagueness and grammar has recently been made by Chierchia (Chierchia 2010) in connection with the distinction between mass and count nouns. This distinction, Chierchia argues, has deep consequences for the grammars of many languages, English among them. The syntactic structure of noun phrases (the various syntactic projections of the category *N*) depends crucially on the properties of the nouns that act in them as 'lexical heads'. In particular the syntactic rules that govern the well-formedness of noun phrases are sensitive to whether or not the noun's extension is built from atoms. For example, the count noun *apple* has an extension that is built from atoms. The *apple*-atoms are the individual apples, and the extension of the plural *apple* consists of all mereological sums that can be formed out of them. Typical examples of mass nouns, such as *mud* or *air*, are widely seen as the exact opposite of this; for instance, what counts as an atomic quantity of mud is completely undefined. But in between these conceptually clear cases there are many nouns – both mass and count – that cannot be characterized straightforwardly in such terms. For example, *furniture* is a 'fake mass noun' in that its extension clearly does have atomic elements (e.g., individual pieces of furniture in a room.) In addition, *segment* is a count noun, although segments typically contain smaller segments as parts, so many elements of the singular noun *segment* aren't atoms and in fact it isn't clear whether any of them are atoms. Other nouns can be used both as mass nouns and as count nouns – *stone, rope, rock, beer, vice* and so on, and it isn't always clear whether their extensions are built from atoms or not. Also, it is sometimes not clear if a noun does have both uses (is there a use of *string* as count noun, with the meaning 'piece of string'; and what about *yarn*?)

Chierchia proposes to deal with these and other problems in the syntax and semantics of noun phrases by assuming that the notion of an '*N*-atom' (where *N* is a noun) is vague and that its vagueness can often be resolved in context. Furthermore, the formation rules for noun phrases built from *N* as lexical head refer to N-atomicity and thus are sensitive to the indeterminacies of this notion, with the result that the notions of well-formedness and Logical Form themselves become indeterminate. Often, Chierchia goes on to argue, the context will make clear which parts of the extension of a noun are to be treated as atoms, so that counting elements of the extension will be possible in that context, and with that evaluation of noun phrases containing numerals, such as *three portions of ice cream*.

Chierchia is certainly right that contextual information of this sort is often needed and that the given context will often supply it. What is not so clear is whether the indeterminacy of atomhood is to be treated as an instance of vagueness and dealt with formally in terms of supervaluation, as Chierchia proposes.[43]

One reason for discounting the indeterminacy of *N*-atom predicates as instances of vagueness is that it does not seem to give rise to sorites problems. But is susceptibility to sorites problems a necessary requirement for ('genuine') vagueness? The community has long been divided on this point. But what depends on the answer to this question? What Chierchia argues for (and plausibly, we think) is that context is often needed to determine what the relevant *N*-atoms are and thus to define the semantic values of various *N*-based noun phrases (among them those that contain numerals). Furthermore, he observes (again, plausibly) that many contexts resolve these indeterminacies only partially, and that contexts are therefore partially ordered in terms of how many unresolved cases each of them settles: $c \leq_N c'$ iff $[\![\mathrm{AT}_N]\!]_c^+ \subseteq [\![\mathrm{AT}_N]\!]_{c'}^+$ and $[\![\mathrm{AT}_N]\!]_c^- \subseteq [\![\mathrm{AT}_N]\!]_{c'}^-$, where $[\![\mathrm{AT}_N]\!]_c^+$ is the positive extension in *c* of the predicate of being an *N*-atom, $[\![\mathrm{AT}_N]\!]_c^-$ the negative extension of that predicate, and likewise for *c*'. As we saw in Section 2, such a structure of partial predicates, whose truth value gaps are partially resolved by contexts and where these contexts form a partial ordering defined by their resolving powers, comes close to reconstructing the supervaluation semantics for vague predicates in its original form. The only missing assumption is that the truth value gap of AT_N gets fully resolved in the limit. Chierchia's discussion of atomicitiy suggests no reasons for resisting this assumption.

The substantive point here is that a semantic characterization in this form of a given predicate *P* does not entail in and of itself that *P* must be governed by tolerance, and thus that it must be liable to the sorites. That makes being vague in a sorites-generating way logically independent

[43] Chierchia's (2010) theory grounds the mass/count distinction in a conceptual distinction between inherently individuable entities and stuff. Rothstein (2014) challenges these grounds and endorses the more radical claim that information about atomicity is **always** added to the interpretation of a noun relative to a context. A count noun is derived by an operation from a mass noun meaning to a set of contextual atoms. The contextual set of atoms is encoded in the type of count nouns.

from the mere fact of having a supervaluation semantics. Is having such a semantics by itself sufficient ground for calling *P* vague? Or should we call predicates with a supervaluation semantics but which are not governed by tolerance by some other name?

With this last question we end. There may not be a single definitive answer to what vagueness is. But why should there be? There are a number of characteristics that we have come to associate with the notion of vagueness in an intuitive and informal manner. A closer and more systematic study will be needed to show how these characteristics are logically related, and thus which combinations of them are logically possible. And a closer investigation into the nature of language and thought can reveal which of those combinations are in fact realized there. In this survey, we have tried to give an impression of the considerable progress that has been made on both these fronts. But the subject is far from exhausted – partly, no doubt, because there are many aspects of vagueness that are still hidden from view; but also because it is so hard to combine the insights gained through work in different disciplines into a single coherent picture.

References

Aarts. Bas. (2004) (eds.), Fuzzy Grammar: A reader. Oxford: Oxford U. Press.

Altshuler Daniel and Schwarzschild Roger (2013) "Correlating cessation with double access", Proceedings of the 19th Amsterdam Colloquium, M. Aloni, M. Franke and F. Roelofsen (eds.), pp. 43-50.

Alxatib, Sam and Francis Jeffry Pelletier (2011). The Psychology of Vagueness: Borderline cases and Contradictions. Mind and Language 26 (3): 287-326.

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

Asher, N., Dever, J. and Pappas, C. (2010). Supervaluationism debugged, *Mind* 118:901–33.

Bale, Alan. Clinton (2011b). Scales and Comparison Classes. Natural Language Semantics 19 (2):169-190.

Barker, Chris (2002) The dynamics of vagueness Linguistics and Philosophy 25: 1-36.

Barker, Chris (2006) *Vaguenei.* In Keith Brown (Ed.) Encyclopedia of Language and Linguistics, 2d edition, volume 13: 294-298. Oxford: Elsevier.

Barker, Chris (2009) Clarity and the grammar of skepticism *Mind and Language* 24.3: 253--273.

Barker, Chris (2013) Negotiating taste, *Inquiry* 56.2-3: 240--257.

Beck, Sigrid (2011). Comparison Constructions. In Claudia Maienborn, Klaus von Heusinger and Paul Portner (Eds.), Semantics: An International Handbook of Natural Language Meaning 2: 1341-1389. Mouton de Gruyter. Berlin.

Beck, Sigrid, Sveta Krasikova, Daniel Fleischer, Remus Gergel, Christiane Savelsberg, John Vanderelst and Elisabeth Villalta (2010). Crosslinguistic Variation in Comparison Constructions. In Jeroen Van Craenenbroeck (Ed.), Linguistic Variation Yearbook 2009: 1–66. John Benjamins Publishing Company. Amsterdam.

Belnap, N.D. (1977), A useful four-valued logic, in J.M. Dunn, G. Epstein (eds.), *Modern Uses of Multiple-Valued Logic*, Dordrecht: Reidel, 8-37.

Booij, Elbert, and Galit W. Sassoon (2013) Big Differences: The standard for `big' as used by adults and children. In Proceedings of the Israel Association of Theoretical Linguistics (IATL) 29. MIT working Papers in Linguistics.

Boolos, George (1991) "Zooming Down the Slippery Slope", *Nous*, 25: 695-706.

Burgess, J. 2001. 'Vagueness, epistemicism, and response-dependence', *Australasian Journal of Philosophy* 79: 507-24.

Burnett, H. (2011). The puzzle(s) of absolute adjectives: On vagueness, comparison and the pragmatic origin of scale structure. In Paperno, D., editor, *UCLA Working Papers in Linguistics: Special Issue on Semantics and Mathematical Linguistics*. University of California, Los Angeles.

Burnett, Heather. (2014a). A Delineation Solution to the Puzzles of Absolute Adjectives. *Linguistics & Philosophy*. 37 (1): 1-39.

Burnett, Heather. (2014b). *Gradability in Natural Language: Logical and Grammatical Foundations*. In press in the series Oxford Studies in Semantics and Pragmatics. Oxford University Press.

Bobalijk, Jonathan David (2012). Universals in Comparative Morphology: Suppletion, superlatives, and the structure of words. The MIT Press. Cambridge, Massachusetts.

Canonica, Caitline 2014 MA thesis, The Hebrew Uiversity, Jerusalem.
Cobreros, Pablo (2011). Supervaluationism and Classical Logic. In Rick Nouwen, Robert van Rooij, Hans-Christian Schmitz and Uli Sauerland (Eds.), Vagueness in Communication. Lecture Notes in Computer Science 6517. Springer. Berlin.
Cobreros, Pablo, Robert van Rooij, Paul Egré and Dave Ripley (2012). Tolerant, Classical, Strict. Journal of Philosophical Logic 41 (2): 347-385.
Cobreros, Pablo, Paul Egré, Dave Ripley and Robert van Rooij (2012). Tolerance and Mixed Consequence in the S'valuationist Setting. Studia Logica 100 (4): 853-876. Special Issue on Logic and Natural Language edited by Nissim Francez and Ian Pratt-Hartmann.
Cobreros, Pablo, Paul Egré, David Ripley, Robert van Rooij. (2014). Vagueness, truth and permissive consequence. In T. Achourioti et al. (Eds.), *Unifying the philosophy of truth*. Springer. (forthcoming).
Constantinescu, Camelia (2011). Gradability in the Nominal Domain. Doctoral Dissertation. Leiden University.
Crespo, Inés and Frank Veltman (2014) Gradable adjectives and subjective experience. Paper draft, ILLC, University of Amsterdam.
Cresswell, Max J. (1976). The Semantics of Degree. In Barbara Partee (Ed.), Montague Grammar: 261-292. Academic Press. New York.
Decock, L., & I. Douven. 'What is Graded Membership?', Noûs 48 (4) (2014), pp. 653-682.
Doetjes, Jenny S., Camelia Constantinescu and Kateřina Součková (2010). A Neo-Kleinian Approach to Comparatives. In Ed Cormany, Satoshi Ito and David Lutz (Eds.), Proceedings of the 19th Semantics and Linguistic Theory Conference: 124–141.
Doetjes, J. (2010). Incommensurability. In Aloni, M., Bastiaanse, H., de Jager, T., and Schulz, K., editors, *Logic, Language and Meaning, 17th Amsterdam Colloquium, Revised Selected Papers*, pages 254–263. Springer, Berlin.
Douven, I., L. Decock, R. Dietz & P. Egré. 'Vagueness: A Conceptual Spaces Approach'. Journal of Philosophical Logic 42 (1) (2013), pp. 137-160.
Dummett, M. 1975. 'Wang's paradox', *Synthese* 30: 301-24; reprinted in his *Truth and Other Enigmas*. Also reprinted in Keefe and Smith (1996).
Egre, Paul, Vincent de Gardelle, and David Ripley (2013) Vagueness and Order Effects in Color Categorization, Journal of Logic, Language and Information 22(4): 391-420.
Fara, Delia Graff. 2000. 'Shifting sands: an interest-relative theory of vagueness', *Philosophical Topics* 28: 45–81.
Fara, Delia Graff. (2010) Scope Confusions and Unsatisfiable Disjuncts: Two Problems for Supervaluationism. In Richard Dietz and Sebastiano Moruzzi (Eds) Cuts and Clouds: Vagueness, Its Truth, and Its Logic, OUP.
Fine, Kit (1975) "Vagueness, truth and logic", *Synthese*, 54: 235-59. Reprinted in *Vagueness: A Reader*, Rosanna Keefe and Peter Smith (eds.), Cambridge: MIT Press, 1996, 119-150.
Fox, Danny (2003). On Logical Form. In Randall Hendricks (Ed.), Minimalist Syntax. Blackwell. Oxford.
Frege, G. 1903. *Grundgesetze der Arithmetik*. Section 56 ff. Translated in P. Geach and M. Black (eds), (1960), *Translations from the Philosophical Writings of Gottlob Frege*, Oxford (2nd edn): Blackwell.
Fox, Danny, and Martin Hackle (2006). The universal density of measurement. *Linguistics and Philosophy* 29:537–586.
Gaifman, Haim (2010). Vagueness, Tolerance and Contextual Logic. Synthese 174 (1): 5–46.
Gamut, L. T. F. (1991). Logic, Language, and Meaning, Volume II: Intensional logic and logical grammar. University of Chicago Press. Chicago.
Gärdenfors, Peter (2004). Conceptual Spaces - The Geometry of Thought. MIT Press. Cambridge, Massachusetts.
Giora, Rachel, Ofer Fein, Jonathan Ganzi, Natalie Alkeslassy Levi and Hadas Sabah (2004). On Negation as Mitigation: The case of negative irony. Discourse Processes 39: 81-100.
Grinsell, Tim (2012) Social choice theory and linguistic vagueness. Salt 22, Chicago.
Hajek, Petr. (2009). Fuzzy Logic. In Edward N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy. Stanford University. Stanford, California.
Hackl, M. (2009). On the grammar and processing of proportional quantifiers: most versus more than half. *Natural Language Semantics*, 17(1):63–98.
Hampton, James (1997). Conceptual Combination. In Koen Lamberts and David Shanks (Eds.), Knowledge, Concepts and Categories. Studies in Cognition: 135-162. The MIT Press. Cambridge, Massachusetts.
Hyde, D. 1997. 'From heaps and gaps to heaps of gluts', *Mind* 106: 440-460.
Hyde, D. 2008. 'Vagueness, logic and ontology', Ashgate, Aldershot.
Jaśkowski, S. 1969 [1948]. 'Propositional calculus for contradictory deductive systems', *Studia Logica* 24: 143-57. Originally published in 1948 in Polish in *Studia Scientarium Torunensis*, Sec. A II: 55-77.

Kadmon, N. and Landman, F. (1993). Any. *Linguistics and Philosophy*, 16(4):353–422.
Kagan, O. & Alexejenko, A. 2010. Degree Modification in Russian Morphology: The Case of the Suffix –ovat. Proceedings of IATL 26, Bar Ilan, Israel.
Kamp, H. 1981. 'The paradox of the heap', in U. Mönnich (ed.), *Aspects of Philosophical Logic*. Dordrecht. Reidel.
Kamp, H. and Partee, B. (1995). Prototype theory and compositionality. *Cognition*, 57(2).
Kamp, Hans. (1975). Two theories about adjectives. In Keenan, E., editor, Formal Semantics of Natural Language. Cambridge University Press, Cambridge, England.
Kamp, Hans (1981). The Paradox of the Heap. In Uwe Mönnich (Ed.), Aspects of Philosophical Logic. Reidel. Dordrecht. Reprinted in von Heusinger, Klaus & ter Meulen, Alice (eds.) 2013. *Meaning and the Dynamics of Interpretation, Selected Papers of Hans Kamp*. Brill: Leiden.
Keefe, R. & Smith, P., eds. 1997. *Vagueness: A Reader*. Cambridge, Massachusetts: MIT Press
Keefe, Rosanna (2000) *Theories of Vagueness*, Cambridge: Cambridge University Press.
Kennedy, C. (1999). *Projecting the adjective: the syntax and semantics of gradability and comparison*. Garland, New York.
Kennedy, C. (2002). The landscape of vagueness. Ms, Northwestern University (available online at http://ocw.mit.edu/courses/linguistics-and-philosophy/24-979-topics-in-semantics-fall-2002/lecture-notes/vagueness.pdf).
Kennedy, C. (2007). Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30(1):1–45.
Kennedy, Chris (2013) Two Sources of Subjectivity: Qualitative Assessment and Dimensional Uncertainty. 2013. *Inquiry* 56.2-3:258-277.
Kennedy, C. and McNally, L. (2005). Scale structure and the semantic typology of gradable predicates. *Language*, 81(2).
Klein, Ewan (1980). A Semantics for Positive and Comparative Adjectives. Linguistics and Philosophy 4: 1–45.
Klein, Ewan (1982). The Interpretation of Adjectival Comparatives. Journal of Linguistics 18: 113-136.
Klein, Ewan (1991). Comparatives. In Arnim von Stechow and Dieter Wunderlich (Eds.), Semantik / Semantics, An International Handbook of Contemporary Research: 673-691. Walter de Gruyter. Berlin.
Krifka, Manfred. 2002. Be brief and vague! And how bidirectional optimality theory allows for verbosity and precision. In David Restle & Dietmar Zaefferer (eds.), Sounds and Systems. Studies in Structure and Change. A Festschrift for Theo Vennemann, 439–458. Berlin: Mouton de Gruyter.
Krifka, Manfred. 2007. Approximate interpretation of number words: A case for strategic communication. In Joost Zwarts Gerlof Bouma, Irene Krämer (ed.), Cognitive Foundations of Interpretation, 111–126. Amsterdam: Mouton de Gruyter.
Krivz, Manuel and Chemla, Emmanuel (2014) Two methods to find truth value gaps and their application to the projection problem of homogeneity. Ms. University of Vienna & LSCP,
Hampton 1997;
Heim I., 2000. Degree operators and scope. Semantics and Linguistic Theory 10, Ithaca, NY, CLC Publications: 40–64.
Labov, William (1973). The Boundaries of Words and Their Meanings. In Charles James N. Bailey and Roger W. Shuy (Eds.), New Ways of Analyzing Variation in English: 340-373. Georgetown University Press. Washington.
Landman, Fred (1991). Structures for Semantics. Kluwer. Dordrecht.
Landman, Fred (2005). An almost (but not quite) Naïve Theory of Measures. Ms. and class notes, Tel Aviv University.
Landman, Fred (2010). Internal and Interval Semantics for CP-comparatives. In Maria Aloni, Harald Bastiaanse, Tikitu de Jager, and Katrin Schulz (Eds.) Proceedings of the Seventeenth Amsterdam Colloquium Conference on Logic, Language and Meaning, Springer Lecture Notes in Computer Science. Springer-Verlag. Berlin, Heidelberg.
Lassiter, D. and Goodman N. (2013). Context, scale structure, and statistics in the interpretation of positive-form adjectives. To appear in Proceedings of SALT 23, Ithaca, NY: CLC Publications.
Lassiter, Dan (2011). Measurement and Modality: The scalar basis of modal semantics. Doctoral Dissertation. Department of Linguistics, New York University.
Lasersohn, Peter. 1999. Pragmatic halos. Language 75(3). 522–551.
Lewis, David, K. (1970). General Semantics. Synthese 22: 18-67. Reprinted in David K. Lewis (1983), Philosophical Papers 1: 189-229. Oxford University Press. New York.
Lewis, David 1979 Scorekeeping in a Language Game. Journal of Philosophical Logic 8 (1979): pp. 339–359.
Lewis, David 1988 vague identity: Evans Misunderstood', Analysis 48 (1988) 128-30.

Lewis, D. 1993. 'Many but almost one', in K. Campbell, J. Bacon and L. Reinhardt (eds.) *Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong*. Cambridge: Cambridge University Press, pp. 23-38.

Luce, Robert Duncan (1956). Semiorders and a Theory of Utility Discrimination. Econometrica 24: 178-191.

McCloskey, Michael and Sam Glucksberg (1978). Natural Categories: Well defined or fuzzy sets? Memory and cognition 6: 462-472.

McConnell-Ginet, S. (1973). *Comparative constructions in English: A syntactic and semantic analysis*. PhD thesis, University of Rochester.

McNally, L. (2011). The relative role of property type and scale structure in explaining the behavior of gradable adjectives. In Nouwen, R., van Rooij, R., Sauerland, U., and Schmitz, H., editors, *ViC 2009 (Papers from the ESSLLI 2009 Workshop on Vagueness in Communication)*. Springer, Berlin.

Morzycki, Marcin (2009). Degree Modification of Gradable Nouns: Size adjectives and adnominal degree morphemes. Natural Language Semantics 17:175–203.

Morzycki, Marcin (2012). The Several Faces of Adnominal Degree Modification. In Jaehoon Choi, E. Alan Hogue, Jeffrey Punske, Deniz Tat, Jessamyn Schertz, and Alex Trueman (Eds.), Proceedings of the 29th West Coast Conference on Formal Linguistics (WCCFL). Cascadilla Proceedings Project. Somerville, Massachusetts.

Morzycki, Marcin. 2014. Modification. Cambridge University Press.

Murphy, Gregory L. (2002). The Big Book of Concepts. The MIT Press. Cambridge, Massachusetts.

Osherson, Daniel N. and Edward E. Smith (1981). On the Adequacy of Prototype Theory as a Theory of Concepts. Cognition 9: 35–58.

Pagin, Peter 2010 vagueness and domain restrictions. In Paul Égré, Nathan Klinedinst. (Eds.) Vagueness and language use. Palgrave .

Panzeri, Francesca & Francesca Foppolo 2011 Can children tell us something about the semantics of adjectives? Pre-proceedings of the 18th Amsterdam Colloquium: 142-151.

Partee, Barbara H. (1987). Noun Phrase Interpretation and Type-shifting Principles. In Jeroen Groenendijk, Dick de Jong and Martin Stokhof (Eds.), Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers: 115–143. Foris Publications. Dordrecht. Reprinted in Paul Portner and Barbara H. Partee (eds.) (2002), Formal Semantics: The Essential Readings: 357–381. Blackwell. Oxford.

Raffman D., (2005) How to understand contextualism about vagueness, Analysis 65 (3), 244-48.

Raffman, Diana. 1994. Vagueness without paradox. The Philosophical Review 103:41–74.

Ravid, Dorit, Elena Tribushinina, Katarina Korecky-Kröll, Aris Xanthos, Marianne Kilani-Schoch, Sabine Laaha, I. Leibovitch-Cohen, B. Nir, Ayhan Aksu Koç, Wolfgang U. Dressler. and Steven Gillis (2010). The First Year of Adjectives: A cross-linguistic study of the emergence of a category. Poster presented at The Child Language Seminar, London, June 24-25, 2010.

Ripley, David (2011). Contradictions at the Borders. In Rick Nouwen, Robert van Rooij, Uli Sauerland and Hans-Christian Schmitz (Eds.), Vagueness in Communication: 169-188. Lecture Notes in Computer Science 6517. Springer. Berlin.

Rips, L. & Turnbull, W. (1980). How big is big? Relative and absolute properties in memory. *Cognition* 8: 145-174.

Rotstein, C. and Winter, Y. (2004). Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics*, 12(3).

Russell, B. 1923. 'Vagueness', *The Australian Journal of Philosophy and Psychology* 1: 84-92. Reprinted in Keefe and Smith (1996).

Sassoon, Galit. 2012. A slightly modified economy principle: stable properties have non stable standards. In Proceedings of the Israel Association of Theoretical Linguistics (IATL) 27. MIT working Papers in Linguistics..

Sassoon, Galit W. 2013. *Vagueness, gradability and typicality, The interpretation of adjectives and nouns*, Leiden: Brill.

Sassoon, Galit W. 2013. A typology of multidimensional adjectives. Journal of Semantics 30(3). 335–380.

Sauerland, U. & Stateva, P. 2010. Two Types of Vagueness. In P. Égré and N. Klinedinst (eds.) Vagueness and Language Use. Palgrave Studies in Pragmatics, Language and Cognition, Palgrave Macmillan.

Serchuk, Phil, Ian Hargreaves and Richard Zach (2011). Vagueness, Logic and Use: Four experimental Studies on vagueness. Mind and Language 26: 540–573.

Shapiro, Stewart (2006) *Vagueness in Context*, Oxford: Oxford University Press.

Siegel, M. E. A. (1979). Measure adjectives in Montague grammar. In Davis, S. and Mithun, M., editors, *Linguistics, Philosophy and Montague Grammar*. University of Texas Press, Austin.

Soames, S. 2002. 'Replies', *Philosophy and Phenomenological Research* 65: 429-52.

Solt, Stephanie. (2014a) Measurement scales in natural language. In press in Language and Linguistics Compass.

Solt, Stephanie. (2014b) On quantification and measurement: the case of most and more than half. Angenommen. Forthcoming in *Language*.
Solt, Stephanie. 2012. Comparison to arbitrary standards. In Proceedings of Sinn und Bedeutung 16. MIT Working Papers in Linguistics, volume 2, 557{570.
Solt, Stephanie, C. Cummins and M. Palmović (2014). The preference for approximation. Unpublished draft.
Schwarzschild, Roger (2005). Measure Phrases as Modifiers of Adjectives. Recherches Linguistiques de Vincennes 35: 207–228.
Sorensen, R.A. 1988. *Blindspots*. Oxford: Clarendon Press.
Sorensen, R. 2001. *Vagueness and Contradiction*. New York: Oxford University Press.
Sorensen, Roy, 2013 *Vagueness, The Stanford Encyclopedia of Philosophy* (Winter 2013 Edition), Edward N. Zalta (ed.) //plato.stanford.edu/archives/win2013/entries/vagueness/.
Stalnaker, Robert (1975). Indicative Conditionals. Philosophia 5. Reprinted in William L. Harper, Robert Stalnaker and Glenn Pearce (Eds.), (1976), Ifs: 193-210. Kluwer. Dordrecht.
Stalnaker, Robert (1978). Assertion. In Peter Cole (Ed.), Syntax and Semantics 9: Pragmatics: 315-332. Academic Press. New York.
Stanley, J. 2003. 'Context, interest-relativity, and the sorites', *Analysis* 63: 269-80.
Stanley, Jason (2003) "Context, interest-relativity, and the sorites", *Analysis*, 63: 269-80.
von Stechow, Arnim. 1984. Comparing semantic theories of comparison. Journal of Semantics 3:1-77.
von Stechow, Arnim. 2009. The temporal adjectives früh(er)/spät(er) and the semantics of the positve. In Quantification, definiteness, and nominalization., ed. Anastasia Giannakidou and Monika Rathert, 214-233. Oxford: Oxford University Press.
Syrett, K. (2007). *Learning about the structure of scales: Adverbial modification and the acquisition of the semantics of gradable adjectives*. PhD thesis, Northwestern University.
Syrett, K., Kennedy, C., and Lidz, J. (2010). Meaning and context in children's understanding of gradable adjectives. *Journal of Semantics*, 27(1):1–35.
Teodorescu, Alexandra (2009). Modification in the Noun Phrase: The Syntax, Semantics, and ragmatics of Adjectives and Superlative. Ph.D Dissertation, The University of Texas at Austin.
Toledo, A. and Sassoon, G. W. 2011. Absolute vs. Relative Adjectives - Variance Within vs. Between Individuals. Proceedings of SALT 21. Rutgers University. New Jersey. MIT working Papers in Linguistics.
Tribushinina, Elena. (2013). Adjective semantics, world knowledge and visual context: Comprehension of size terms by two- to seven-year-old Dutch-speaking children. *Journal of Psycholinguistic Research* 42(3): 205-225.
Tye, Michael (1994) Sorites paradoxes and the semantics of vagueness. Philosophical Perspectives 8:189-206.
Unger, P. (1975). *Ignorance: A Case for Scepticism*. Clarendon Press, Oxford.
Unger, Peter (1979) "There are no ordinary things", *Synthese*, 4: 117-54.
Unger, Peter (1980) "The Problem of the Many" *Midwest Studies in Philosophy* 5: 411-67.
van Benthem, J. (1982). Later than late: on the logical origin of the temporal order. In *Pacific Philosophical Quarterly*, volume 63, pages 193–203.
van Deemter, Kees (2010). Not Exactly: In praise of vagueness. Oxford University Press. Oxford.
van Fraassen, Bas C. (1969). Presuppositions, Supervaluations and Free Logic. In K. Lambert (Ed.), The Logical Way of Doing Things: 67-91. Yale University Press. New Haven.
van Rooij, Robert (2011a). Measurement and Interadjective Comparisons. Journal of Semantics 28 (3): 335-358.
van Rooij, Robert (2011b). Implicit versus Explicit Comparatives. In Paul Egré and Nathan Klinedinst (Eds.), Vagueness and Language Use: 51–72. Palgrave Macmillan. New York.
van Rooij, Robert (2011c). Vagueness and Linguistics. In Giuseppina Ronzitti (Ed.), Vagueness: A guide: 123–170. Springer. London, New York.
van Rooij, Robert, Pablo Cobreros, Paul Egré and Dave Ripley (2012). Tolerant, Classical, Strict. Journal of Philosophical Logic 41 (2): 347-385.
Veltman, Frank (1984). Data Semantics. In Jeroen Groenendijk, Theo Janssen and Martin Stokhof (Eds.), Truth, Interpretation and Information. Proceedings of the Third Amsterdam Colloquium. Foris Publications. Dordrecht.
Veltman, Frank (1996). Defaults in Update Semantics. Journal of Philosophical Logic 25: 221-261.
Venneman, Theo and Renate Bartsch (1972). Semantic Structures: 47-86. Anthenäum Verlag. Frankfurt.
von Fintel, Kai (1994). Restrictions on Quantifier Domains. Doctoral Dissertation. University of Massachusetts. Amherst.
Wellwood, A., V. Hacquard, and R. Pancheva (2012). "Measuring and Comparing Individuals and Events" *Journal of Semantics* 29, 207-228.
Williamson, T. 1994. *Vagueness*. London: Routledge. Chapters 7-8.

Williamson, T. (1999) On the structure of higher-order vagueness, *Mind* 108 (429), 127-143.
Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press
Wright, C. 1975. 'On the coherence of vague predicates', *Synthese* 30: 325-65.
Yen, John (1999). Fuzzy Logic—A Modern Perspective. Transactions on Knowledge and Data Engineering 11 (1): 153-165.
Zadeh, L. 1975. 'Fuzzy logic and approximate reasoning', *Synthese* 30: 407-428.
Zwicky, A. M. and J. M. Sadok: 1975, 'Ambiguity Tests and How to Fail Them', Syntax and Semantics 4, 1-36.