Vagueness and Domain Restriction*

Peter Pagin

December 11, 2010, 14:13

Abstract

This paper develops an idea of saving ordinary uses of vague predicates from the Sorites by means of *domain restriction*. A tolerance level for a predicate, along a dimension, is a difference with respect to which the predicate is semantically insensitive. A central gap for the predicate+dimension in a domain is a segment of an associated scale, larger than this difference, where no object in the domain has a measure, and such that the extension of the predicate has measures on one side of the gap and the anti-extension on the other. The domain restriction imposes a central gap.

1. Tolerance principles

In the introduction to their vagueness reader, Rosanna Keefe and Peter Smith classified accounts of vagueness with respect to how they handle the sorites paradox.

---

*The ideas in this paper were presented at the conference *Vagueness and Language Use*, in Paris, April 2008. I am grateful to several of the participants for comments, including Paul Egré, Robert van Rooij, Agustin Rayo, and Stephen Schiffer. I am grateful as for helpful comments from an anonymous referee, as well as from the editors. And, as always I am indebted to Kathrin Glüer-Pagin for discussion of the problems over the years.

Accounts of vagueness similar to the present one, giving an important role to *gaps* in the domain, have been proposed by Ruth Manor (2006), Haim Gaifman (2002), Robert van Rooij (this volume), and Mario Gómez-Torrente (2010). For brief comparisons between those accounts and the present proposal, see footnote 14.

I first came up with these ideas when in 1997 reading Williamson 1994. I presented them in two international conferences in 1998 (Putnam conference in Karlovy Vary, *Meaning and Interpretation* conference in Stockholm). I took up the theme again in a grant application in 2004, and presented the basics of the account in a vagueness workshop in St Andrews in 2006 (published as Pagin 2010). I was told about Manor's work by Tim Williamson after he learned about mine, while Gaifman's work was unknown to me until after completing the first draft of this paper, brought to my attention by an anonymous referee.

The work was sponsored by the Swedish Research Council (VR) grant *Vagueness and Context Factors.*

The sorites paradox is set out in the standard way with reference to a sorites *sequence s* of objects $s_1, \ldots, s_n$ and an associated vague predicate F. In S, there is a very small and seemingly negligible difference between any two adjacent elements s_i and s_{i+1} with respect to the dimension that is relevant to satisfying F (for instance, if F is '. . . is tall', then the dimension is height). This suggests that if s_i satisfies F, then so does s_{i+1}. Since S is a sorites sequence for F it is also stipulated that s_1 satisfies F and that s_n does not. Let t_i denote s_i, $1 \leq i \leq n$. Then the sorites argument is set up as

(1)
1 $F(t_1)$
2 $\forall i(F(t_i) \rightarrow F(t_{i+1}))$

---

3 $F(t_n)$

where premise 1 is apparently true, premise 2, the inductive premise is apparently true, the argument is apparently valid, and the conclusion 3 is apparently false. Keefe and Smith then say

> The standard alternatives for coping with the sorites paradox are
>
> a) reject the validity of the argument
>
> b) question or reject the (strict) truth of the inductive premise
>
> c) reject the truth of the minor premise or the falsity of the conclusion
>
> d) accept the whole reasoning and conclude that the vague predicate is incoherent (Keefe and Smith 1996, 10).

I think it is fair to say alternative b) has been the most popular one in the literature of recent decades, but this is not the place to review the alternatives.

In Pagin 2010 I presented an alternative strategy for handling the sorites paradox.[1] This strategy is not stated with respect to induction over a sorites sequence but with respect to *general* inductive principles. Roughly, the strategy consists in accepting alternative a), *reject validity*, for *ordinary* uses of vague vocabulary, and alternative d), *incoherence*, for "extreme" uses. The use of a vague predicate in a sorites sequence, on this view, is an extreme use. Because of this, the proposed

[1] Concerning similar accounts and the question of originality, see the note at the end of section 6.

strategy is not an attempt to *solve* the sorites paradox. Rather, the contradiction is accepted. The main idea is to provide a coherent semantics for ordinary uses of vague vocabulary, *with* (and adequate formulation of) the inductive principle in place.

The inductive principle, stated for the predicate 'tall' and with respect to a difference of 1 mm, can be stated as

(2) If a man of $n+1$ mm is tall, then a man of n mm is tall

This principle is an inductive counterpart to the statement that

(3) 1 mm cannot make the difference between a man's being tall and a man's not being tall

(3) expresses the intuition that 'tall' is insensitive to small differences, together with the intuition that 1 mm is small enough. The property of being insensitive to small differences is what Crispin Wright has called *tolerance*. Wright says

> What is involved in treating these examples as genuinely paradoxical is a certain *tolerance* in the concepts which they respectively involve, a notion of a degree of change too small to make any difference, as it were. ...Then F is *tolerant* with respect to ϕ if there is also some positive degree of change in respect of ϕ insufficient ever to affect the justice with which F applies to a particular case (Wright 1976, 156).

For reasons to be given later, I prefer to speak of linguistic *expressions* as being vague or tolerant, rather than concepts. I shall refer to inductive principles such as (2) as *tolerance principles*. I shall refer to the particular difference given in the principle as the *tolerance level*.

Tolerance intuitions regarding vague expressions are widespread. Most accounts of vagueness involve a rejection of tolerance. Usually, they are coupled with explanations of why we have those intuitions despite the falsity of the tolerance principles.[2] In some cases, tolerance is nominally "accepted", but then in the

[2] For instance, Tim Williamson (1994) offers his elegant margin-of-error principle as justification for explaining why we cannot have knowledge of the application of vague predicates in borderline

form of a principle for *judging*. For instance, Stewart Shapiro says

> Suppose that two objects, a, a' differ only marginally in the relevant respect (on which P is tolerant). Then if one competently judges a to have P, then one cannot competently judge a' in any other manner (Shapiro 2006, 8).

and refers to this as a principle of tolerance. Shapiro does not, however, accept tolerance in the present sense, i.e. *semantic* tolerance, as stated e.g. in (2).[3]

The leading idea of the present approach is to accept the tolerance intuitions, and to devise a semantics that nonetheless gives a coherent representation of ordinary use of vague vocabulary. Two simple observations lie at the bottom of the approach. First, an informally stated tolerance principle such as (2), is ambiguous between two different regimentations into first-order logic:

(T1) $\forall y(\text{Man}(y) \;\&\; H(y) \geq n+1 \;\rightarrow\; \text{Tall}(y)) \;\rightarrow$
$\forall y(\text{Man}(y) \;\&\; H(y) \geq n \;\rightarrow\; \text{Tall}(y))$

(T2′) $\forall x \forall y((\text{Man}(x) \;\&\; \text{Man}(y) \;\&\; H(x) \leq n+1 \;\&\; H(y) \geq n) \;\rightarrow$
$(\text{Tall}(x) \;\rightarrow\; \text{Tall}(y)))$

where $H(x)$ is the height of x (where H is the relevant measure function for the predicate 'tall'; see below). These principles are non-equivalent. The second is equivalent to

(T2) $\exists y(\text{Man}(y) \;\&\; H(y) \leq n+1 \;\&\; \text{Tall}(y)) \;\rightarrow\; \forall y(\text{Man}(y) \;\&\; H(y) \geq n \;\rightarrow\; \text{Tall}(y))$

The sharpest difference between (T1) and (T2) concerns domains with *central gaps*.[4]

2. Measure functions and central gaps

To define the concept of a central gap we first need the concepts of an *associated weak order*[5] and a *measure function*:

cases. But this does not by itself account for the fact that speakers *believe* in the truth of those principles.

[3] Tolerance is accepted in by Elia Zardini (2008). Zardini avoids the sorites paradox by a restriction on logical consequence.

[4] (T1) and (T2) are characterized model-theoretically in Pagin 2010.

[5] A weak order $\lesssim$ over a domain D is a binary relation over D satisfying the conditions

(AWO) Let F be a one-place predicate over a domain D. We call $\lesssim_F$ an F-associated weak order iff

i) $\lesssim_F$ is a weak order over D

ii) For all $d, d' \in D$, if F is true of d and $d \lesssim_F d'$, then F is true of d'.

By condition ii), the so-called *penumbral* connections of F must hold with respect to the F-associated weak order. Next, the definition of a *measure function*:

(AMF) Let F be a one-place predicate over a domain D. $\mathcal{H}_F$ is a *measure function* for F iff

i) $\mathcal{H}_F$ is a function from D to Re (the set of real numbers)

ii) there is an F-associated weak order $\lesssim_F$ over D such that for all $d, d' \in D$,

$$\mathcal{H}_F(d) \leq \mathcal{H}_F(d') \quad \textit{iff} \quad d \lesssim_F d'.$$

Here $\langle Re, \leq \rangle$ is the relational structure of the Reals with the *less than or equal to* relation. Thus, the measure function is a homomorphism from $\langle D, \lesssim_F \rangle$ to $\langle Re, \leq \rangle$. The existence of an associated weak order is usually assumed for vague predicates, and for any finite or countable weak order, a homomorphism into $\langle Re, \leq \rangle$ exists (cf. Krantz et al. 1971, pp. 15-17, 39).

However, if we only have the condition that the measure function be a homomorphism, then any order-preserving function between the two relational structures qualifies as a measure function. That is, the measure function is unique only up to order-preserving transformations, which is to say that we have a so-called *ordinal scale*. An ordinal scale may be enough for the sorites susceptibility of tolerant predicates in case we only consider measure functions into N, the set of natural numbers. This may be the case e.g. when we compare entities with respect to baldness (number of hairs) or heap-hood (number of grains), where only whole numbers matter for the comparison. For then the tolerance of 'bald' amounts to the fact that if anyone is bald, then there is or can be someone who has more hair

i) For all $d, d' \in D$, $d \lesssim d'$ or $d' \lesssim d$ (connectedness)
ii) For $d, d', d'' \in D$, if $d \lesssim d'$ and $d' \lesssim d''$, then $d \lesssim d''$ (transitivity)

and still is bald. Having more hair consists in having at least one more individual hair. Then, by repeating the tolerance step any finite boundary of baldness can be crossed, and we will have a sorites contradiction.

This does not work where the order is dense or continuous, as with height. The *taller than* order (over possible objects) is dense, for between any two objects x, y such that x is strictly taller than y, there is a third object z such that z is strictly taller than y and strictly shorter than x. If in this case tolerance does not amount to more than what can be captured by an ordinal scale, it only consists in the circumstance that for any tall object there is or can be an object that is strictly shorter but still tall. But it is then possible that for a given tall person A of 180 cm there are infinitely many shorter persons of different height that are tall, although no one of a height less than 179 cm is tall. This is so simply because there are infinitely many rational and hence also real numbers between 180 and 179. So in this case, tolerance in the weak, ordinal, sense does not induce a sorites contradiction.

What we need for tolerance in the normal sense expressed by principles such as (2) is the idea of a margin of a certain fixed size, such as 1 mm, such that no two objects whose difference is within that margin can differ with respect to satisfying some particular predicate.[6] An ordinal scale does not offer that. From assignments on an ordinal scale we cannot get the information that the difference between two objects is at least as great as the difference between two other objects. We need an *interval scale*, i.e. a measure function that is unique *up to linear transformations*. On an interval scale the ratios between differences are fixed. If

$$\frac{\mathcal{H}_F(d) - \mathcal{H}_F(d')}{\mathcal{H}_F(a) - \mathcal{H}_F(a')} = r$$

and $\mathcal{H}'$ is an admissible (hence linear) transformation of $\mathcal{H}$ on an interval scale, then also

$$\frac{\mathcal{H}'_F(d) - \mathcal{H}'_F(d')}{\mathcal{H}'_F(a) - \mathcal{H}'_F(a')} = r.$$

Specifically, if $r = 1$, the differences are identical. So interval size identity does not vary between admissible functions. In order to construct an interval scale mea-

[6]Strictly speaking, the sorites contradiction arises even if the tolerance margin d_i between s_i and s_{i+1} depends on i and is monotone decreasing ($d_{i+1} < d_i$), as long as the series $d_1 + d_2 + \ldots$ diverges.

sure function, one needs a procedure for setting intervals as equal (cf. Krantz et al. 1971, chapter 1). The most direct way of doing so is given by the existence of a *concatenation operation* $\circ$, i.e. an operation of combining two objects such that the combination can be compared with other objects or combinations. In the case of weight we can weigh two material objects together, and in the case of length we can put them end-to-end. An operation $\circ$ on a domain D qualifies as a concatenation operation iff for all $d, d' \in D$

$$d \circ d' \succ d, d'$$

(where $\succ$ is *strictly greater*). The further condition on a measure function $\mathscr{H}$ is then that

$$\mathscr{H}(d \circ d') = \mathscr{H}(d) + \mathscr{H}(d').$$

With a concatenation operation, we have *extensive* measurement.[7]

Given the relevant kind of measure functions, i.e. measure functions that satisfy conditions at least as strict as interval scale uniqueness, we can define the concept of a *central gap*. Let $E(F)$ be the extension of F:

(CG) A domain D has a central gap, with respect to a tolerant predicate F, a measure function $\mathscr{H}_F$ and a tolerance level $\mathscr{T}_F > 0$ for F just in case there is a bounded real number interval (i, j) such that

1) $\forall d \in D(d \in E(F) \leftrightarrow \mathscr{H}_F(d) > i)$
2) $\forall d \in D(d \notin E(F) \leftrightarrow \mathscr{H}_F(d) < j)$
3) $i - j \geq \mathscr{T}_F$

For the example of height and ‘tall’ and some domain D, this amounts to saying that a central gap in D is an interval (i, j) in, say, millimeters, such that everyone in D is tall iff he is more than i mm, and everyone in D fails to be tall iff he is less than j mm in height, and the difference between i and j is at least as great as the tolerance level. For instance, with a tolerance level of 1 mm we might have

[7] One can then create a so-called *standard sequence* by choosing an object e as having a unit measure, $\mathscr{H}(e) = 1$, and comparing arbitrary objects with n-ary concatenations of exact replicas of e (cf. Krantz et al. 1971, chapter 1).

a domain with a central gap of 20 mm; everyone is e.g. either taller than 180 cm or shorter than 178 cm.[8] Note that, despite what the term suggests, a central gap need not be "central": the definition allows that the set of Fs in D is either D or $\varnothing$, in which case there is no sorites sequence anyway. Let's say that a central gap for F in D is *significant* just in case there are both Fs and non-Fs in D.

The second simple observation is that in a domain with a significant central gap the (T1) principle is false and the (T2) principle is true but does *not* induce the sorites paradox. To verify the first claim, look at the lower end point of the gap, assume that there is a man in the domain that is 1779 mm and (regarded as) not tall, and consider an instance of (T1):

(4) $$\forall y((\text{Man}(y) \,\&\, H(y) \geq 1780) \rightarrow \text{Tall}(y)) \rightarrow$$
$$\forall y((\text{Man}(y) \,\&\, H(y) \geq 1779) \rightarrow \text{Tall}(y))$$

Clearly, because of the gap, the antecedent is true: everyone in the domain that has a height of 1780 mm or higher also has a height that is more than 1800 mm, and hence he is tall. The consequent is false, however, since there is a man of 1779 mm who is not tall. Hence the conditional is false, and therefore also the principle (T1).

For the second claim, first note that at the lower end point of the gap and above, (T2) instances will have true consequents:

(5) $$\exists y(\text{Man}(y) \,\&\, H(y) \leq 1781 \,\&\, \text{Tall}(y)) \rightarrow$$
$$\forall y(\text{Man}(y) \,\&\, H(y) \geq 1780 \rightarrow \text{Tall}(y))$$

This is because everyone who has a height of at least 1780 mm also has a height of more than 1800 mm, and hence is tall. Second, note that at the upper end point of the gap and below, (T2) instances will have false antecedents, since there is no man in the domain of height at most 1800 mm who is tall. Hence, every instance has a false antecedent or a true consequent. Hence, (T2) is true.

[8] More realistically, the right hand side of condition (CG2) should be weakened: either the measure of d is less than j, or the measure function *is not defined* for d, since not everything has a length. I shall ignore this complication in the rest of the paper.

For the third claim, observe that in order to detach a false consequent

$$\forall y(\text{Man}(y) \mathbin{\&} H(y) \geq n \;\rightarrow\; \text{Tall}(y))$$

we need a true antecedent

$$\exists y(\text{Man}(y) \mathbin{\&} H(y) \leq n+1 \mathbin{\&} \text{Tall}(y)).$$

But if the consequent is false, then $n < 1780$, and because of the gap, there is no true antecedent with $n \leq 1800$. Hence, even if (T2) is taken as an axiom, in a domain with a significant central gap no instance will occur in a modus ponens inference with a true minor premise and false conclusion. With this form of tolerance, gappy domains are safe from the sorites.

Let's say that a domain is *relaxed* with respect to a vague predicate F, or F-relaxed, iff it has a central gap, and that it is *tight* with respect to F, or F-tight, if it does not. If the domain is F-tight, both (T1) and (T2) induce the sorites paradox for F. These are domains that contain sorites sequences. There are many predicates such that the domain of the real world is tight. For instance, given that no man has much more than 200.000 hairs on his head, it is highly probable, given the population of men on Earth, that for each $i < 200.000$, there is a man with i number of hairs on his head. If one hair cannot make the difference between being bald and not being bald, then it is probable that the domain of men on Earth is tight with respect to the predicate *bald*. How can we then make the use of 'bald' in normal circumstances safe from the sorites if tolerance in the form of (T2) is to be accepted?

The proposal in Pagin 2010 was to apply the two familiar ideas of *standard of comparison* and *quantifier domain restriction* in the interpretation of natural language utterances and implement the strategy of central gaps by means of them. This proposal is in part further motivated by a plausible view of the psychology of applying vague predicates. When a speaker says something like

(6) Sam is tall

she typically tacitly has in mind some clearly positive range of tallness and some clearly negative range of tallness, and places Sam in the upper category. She *disregards* the problematic intermediate range of heights, since there is no need to take a stand on heights in that range in order to properly classify Sam.

To be in "the intermediate range" is not the same as being a *borderline* case in any strict sense. If x is a borderline case of tallness, in some context, then x is in the intermediate range of tallness for that context, but the converse does not hold. What I call "the intermediate range" is a device used to model the psychology of a language user. In what follows I shall *not* assume that a context of utterance uniquely *fixes* an intermediate range, only that there are contextually *admissible* ways of using a central gap semantics for representing that feature of speaker psychology.

Such a semantics can be said to provide an account of vagueness provided it gives reasonable interpretations of the vast majority of ordinary uses of vague expressions, even though it does not in the traditional sense offer a solution to the sorites paradox. Rather, it shows how vague terms can be used in ordinary discourse without generating the paradox. But if too many ordinary uses turn out to be incoherent on the proposed semantics, or implausibly interpreted, then (unless there are better implementations of the basic idea), the approach itself will have failed.

I now turn to the task of setting out and developing this idea.

3. Comparisons: classes, standards and gaps

The present account will rely heavily on an appeal to *context dependence*. Many different ideas about context dependence have been used in accounts of adjectives in general, and vague adjectives in particular. One well-established idea in this area is that of a *comparison class*. It has some similarity to the present proposal, since one could think of implementing the central gap strategy by means of comparison classes.

The idea of a comparison class comes out naturally when considering utterances of (6) in different contexts. Uttered in the context where the conversation concerns students in a high school class, the statement may be true, while uttered in a different context where the conversation concerns the members of a basketball team, the statement made *there* may yet be false. This comes out in qualifications with *for*-type prepositional phrases:

(7) Sam is tall for a high school student

(8) Sam is tall for a high school basketball player

In (7) one (typically) uses the set of high school students as comparison class for heights, and this is made explicit in the sentence by means of the PP. Analogously with (8).

The same comparison classes can also be used implicitly with just (6). This idea was developed by Ewan Klein (1980). Klein's basic idea was to have a function $\mathscr{U}$ that for each context of utterance c picked out a subset $\mathscr{U}(c) \subseteq U$ of the general domain U of individuals (Klein 1980, 14). (6) is then true in a context c just in case John counts as tall with respect to the comparison class $\mathscr{U}(c)$.

There are several problems with this simple model. For one thing, it does not allow for variation of comparison class within quantified sentences. Syntax may require contextual updates without any corresponding change in the *external* context.[9]

A second problem is that the use of complex predicates may involve *different* standards or comparison classes, each related to simple constituent:

(9) Sam is neither tall nor strong

[9] There are quantified sentences like

(i) Everyone in my family is tall

which Christopher Kennedy (1999; 2007, 8) glosses as

(i′) For every x in my family, x has a height greater than the norm for someone like x, where the relevant kind of similarity (same age, same sex, etc.) is contextually determined.

If there is only one comparison class determined by context, then this reading where comparison class or standard may depend on the individual to which tallness is ascribed, is not available.

It is not so clear how to handle this phenomenon. I am inclined to think that in the case of (i) there is tacit pragmatic enrichment to 'tall', which in fact makes the property ascribed *different* in each universal instantiation (like e.g. *tall for her age*), rather than the same, and this calls for a treatment that is anyway orthogonal to issues discussed here.

under a reading where a basketball team is intended to provide the comparison class for height and a wrestling team the comparison class for strength.

A third problem, as stressed by Delia Graff Fara (2000) and by Kennedy (2007), is that a comparison class does not automatically give you a *standard* for determining truth value. For instance, to be tall with respect to some comparison class does not necessarily mean to be above the *average* height of members of the class. It seems then, as Kennedy concludes (2007, 9-10), that we need a contextually determined standard of comparison anyway. It may seem as if the comparison class that is salient in the context suggests some standard, since both the comparison class and the standard are made salient in the context. However, if we do have a standard, then we have what is needed for assessing applications of the context-dependent predicate and do not need a comparison class for this purpose over and above the standard.

Contextually determined standards are usually taken to impose sharp boundaries. The existence of sharp boundaries *prima facie* runs counter to intuitions about vague predicates and about applications of them in ordinary practice. This is compensated for in standard contextual theories by letting the context shift and its associated boundary *move*, as a function of judgments made, of attention shifts, or of assertions made in the conversation. One problem with this strategy is that boundaries are taken to move unbeknownst to or even against the intentions of the speakers.[10] Postulating context shifts under such conditions in order to save consistency may seem *ad hoc*. It is e.g. not independently clear why mere change of attention should change the extension of a vague predicate.

In Kennedy's account, the standard of comparison is contextually determined "in such a way as to ensure that the objects that the positive form is true of 'stand out' in the context of utterance, relative to the kind of measurement that the adjective encodes" (Kennedy 2007, 17). This is to ensure by a semantic-pragmatic mechanism that assertions made are safe, i.e. do not apply the vague predicate very close to the boundary. As far as I understand, predications close to the bound-

[10] Shapiro (2006, 35) calls this "the Heraclitus problem" and objects (rightly) that rapid changes in extension of predicates reduces the usefulness of classical logic (I guess it would reduce the usefulness of several other logics as well). Shapiro's own conversational score account is meant to minimize this problem with contextual accounts.

ary will then have the effect of shifting the standard, therefore moving the boundary and hence (perhaps) changing the extension.

One of the goals of the present strategy is that the shift of contexts be controlled by syntax and pragmatics in ordinary ways. Syntax plus semantics may require that the context c at a particular time t of a conversation is a function of the immediately preceding context c_0, e.g. by letting the domain of c be a restriction of the domain of c_0. I shall not postulate context change in order to save consistency. On the contrary, controlled and conservative context shifts may, as we shall see, be a way in which inconsistency is produced (in this framework, context shift is not a method of salvation, but a possible road to damnation). This is not quite the whole story, however, since some context shifts may liberate the speaker from constraints induced by earlier contexts.

In the present proposal, central gaps will serve two interrelated functions. One of these functions will be a counterpart to the function of a standard of comparison in other accounts, such as Kennedy's. With respect to a predicate F and a context c, the central gap will partition the contextual domain so that the extension of F at c consists of the those objects in the contextual domain whose relevant measures are above the gap, and so that the anti-extension consists of those objects whose measures are below the gap. In this role the central gap will serve to fix the truth values of atomic open or closed sentences Ft.

4. Precision: standards, tolerance levels and gaps

David Lewis (1979) introduced the idea of a *standard of precision* as a truth-determining factor of the context of utterance. The standard of precision settles how lax the application of predicate is allowed to be, i.e. how much an object can deviate from strictly having the corresponding property and still count as satisfying the predicate in the context. For instance, J.L. Austin's sentence

(10) France is hexagonal

is true with respect to a low standard of precision and false with respect to any high standard of precision (Austin 1975, 143; Lewis 1979, 352).

It would seem natural to simply equate Lewis's idea of a standard of precision with the graded conception of tolerance, i.e. with the idea of a *level* of tolerance. In particular, this would seem natural, since a standard of *maximal* precision corresponds to a *zero* tolerance level: there is a sharp boundary for the predicate and no deviations from that boundary are allowed.

However, on closer inspection, the two ideas are rather different, at least in their immediate application. That the standard of precision can vary, and that it can be low, does not entail tolerance, and is even inconsistent with tolerance (at least of the (T1) variety). On Lewis's account, a non-maximal standard of precision still determines a sharp boundary of application of the predicate, only one that allows a wider extension than does the maximal precision standard. With respect to some standard, France is *just* hexagonal-like *enough* to be hexagonal by that standard; anything less hexagonal-like, however little, would not be.

The idea of a standard of precision is analogous to the idea of a *margin of error* in measurement: it tells you the maximal distance between the measured value of some quantity and the true value of that quantity. It does not tell you that it would be correct to assign to the quantity any other value within the margin of error. So in one respect, the idea of tolerance is the opposite: it is analogous rather to an idea of a *margin of accuracy*. With tolerance, it is a necessary condition of an object o to satisfy some predicate that any object o' that differs from it along the relevant dimension, up to a certain maximal distance, satisfies it too. Since this necessary condition applies again to o', the sorites can be generated.[11]

Because of this, there is a sense in which tolerance is *reverse precision*. Precise predicates have zero tolerance, maximal reverse precision, while vague predicates have non-maximal reverse precision, non-zero tolerance, with a value that depends on the context. So while tolerance level, as a contextually determined value, cannot be fully assimilated to the independently motivated idea of a standard of precision, it is still clearly related to it.

[11] In Williamson 1994, there is in this sense a margin for accuracy of *knowledge*: a subject can know that an object o has a property F only if any object o' differing from o in a certain maximal way, has the property F as well. This margin of accuracy is in turn induced by a margin of error of the method of measuring employed (determining the number of leaves on a tree by looking at it from a distance).

5. Domain restriction

The second function of central gaps in the present account is to effect a domain restriction for quantifiers of the object language. The need for quantifier domain restriction is usually introduced by means of examples such as

(11) Everyone left at midnight

where what is meant is that everyone who was present e.g. at some particular party, left (the location of the party) at midnight. The domain of 'everyone' in (11) is therefore restricted to the set of those taking part in the event. The need to take account of these phenomena in formal semantics has been acknowledged for a long time. A seminal early contribution is Westerståhl 1985.

As suggested by Westerståhl, a domain restriction is effected by means of a context set that is intersected with the antecedently given domain. In case we have binary quantification, of the form *Q*AB, where *Q* is a determiner, the context set is intersected with the set denoted by the first argument, A. To take an example from Jason Stanley and Zoltan Zsabó (2000), with a context set M, the sentence

(12) Every man runs

is given the truth conditions

(12′) T('Every man runs', *c*) iff $(\mathbf{Every}(\mathbf{man}\cap \mathrm{M}))(\mathbf{run})$

Stanley and Zsabó (2000) argue that the phenomenon of domain restriction should be given a semantic rather than a syntactic or pragmatic account. As they define the terms, on a *semantic* account, the context set is given as a contextually determined value to variables in logical form. They further argue that this variable should belong in the nominal that constitutes the first argument to the determiner ('every' in (12)).

The present account will be semantic in the more general sense that context dependence involves assigning values to context variables provided in the semantic theory. It is often immaterial, however, whether context dependence is represented in the syntax, and therefore also immaterial whether the account is semantic in the more narrow sense of Stanley and Zsabó.[12]

Here I shall follow the tradition in representing a contextually determined domain restriction as an intersection with the first argument of the determiner denotation. In order to implement the domain restriction of the central gap account, there will be two further restrictions, assuming that both the first and the second argument of the determiner is vague. In a particular context c, a central gap will be associated with each predicate in the lexicon that has been used in the context. Let us assume that each vague lexical item is associated with a measure function. Where F is a predicate, let $\nabla(c, F)$ be the *F domain cut*, the set of objects whose measures are within the central gap associated with F in c. We take it to return the empty set in case the predicate in question is not vague. Let $\nabla(c, \overline{F})$ be

$$\bigcup_{i=1}^{n} \nabla(c, F_i),$$

the union of domain cuts for predicates $F_i, \ldots, F_n$ in context c. For two atomic predicates used together, the union is represented in Figure 12 as comprising two intersecting bands in the plane.[13] In the restricted domain, no object has a pair of measures in either band. With respect to example (12) we will then have a domain that contains (at most) $(\mathbf{man} \cap M) - (\nabla(c, \mathit{man}) \cup \nabla(c, \mathit{runs}))$.

With such extra domain restrictions, we will have what is needed to avoid the sorites paradox in relation to the (T2) principle. This is the first function of the central gap, which is achieved by means of the second: let the extension contain objects with values above the gap, the anti-extension objects with values below the gap, and let there be no objects in the domain with values in between.

6. Outline of the semantics

The semantics that I shall propose has the following properties: a *gap function* $\mathscr{G}$ has the role of an interpretation function. It is defined for a predicate F in a domain D iff D is F-relaxed or allows a *restriction* that makes it F-relaxed. When this

[12]This is exemplified in Pagin 2005.

[13]This can also be regarded as the two-dimensional gap for complex predicates such as 'is tall and fat' and 'is tall or fat'.

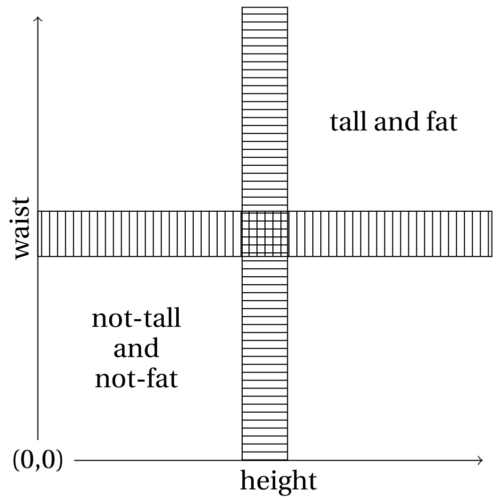

Figure 1: Union of gaps.

condition is met for a context c, the interpretation of speakers in c gives a classical (bivalent) semantics where the (T2) principle is true but there are no sharp boundaries for the predicates.

It will be assumed that in each context c, for each vague predicate F that is used in c, a central gap is determined for F. This is an idealization. In any actual occasion of speech, it will be left undetermined exactly what the tolerance level is and where the gap is. It would be more realistic to speak of what is *admissible* on a certain occasion. This is all the more realistic since the contextual parameters will have to depend to some extent on what is uttered on the occasion.

It will be required that a full semantics is given in each context for the full *fragment* of a language that is used in the context, but not for linguistic material outside that fragment. This means in particular that any referring singular term that is used in c must have referent in the quantifier domain of c. For maintaining bivalence, the relevant measure of that term for the predicate F cannot fall in the central gap for F at c. Hence, the central gap is required not to include the measure of that referent. So the location of the gap depends on the terms used in the context. This will be a requirement on any function $\mathcal{G}$ that maps contexts and predicates on central gaps. Nonetheless, it is formally an unnecessary complication to consider a *family* of such gap functions and assume that one of these functions is selected in each context. I shall therefore operate with just one such function.

This function must to some extent be arbitrary (a choice function), but must also be subject to a number of restrictions, ranked according to importance.

(GAP) i) The central gap must be selected so that the full fragment of language used in the context, including pragmatically determined contextual updates, is taken account of.

ii) The size of the central gap must be at least equal to the tolerance level.

iii) The location of the central gap must be selected so as to ensure consistency.

iv) The location of the central gap must selected so that what the speaker says comes out as true, to the extent this is possible and reasonable.

The order of these selection principles is meant to reflect the order of importance. Indeed, (GAPi) cannot be compromised, but this has the consequence that the other three may not be jointly satisfiable, given collateral facts. When they are not, it is not always clear e.g. that consistency of the speaker should take precedence over the truth of individual judgments. But this is a matter of further investigation.

It is psychologically plausible that the gap is greater, and even considerably greater, than the tolerance level, but for the purpose of setting out the semantics, it is more elegant to reduce the two contextual parameters, tolerance level and gap size, to a single parameter. We can let the gap size *equal* the tolerance level. Then, the requirement (GAPii) will be automatically met.

We can further make the natural stipulation that the *standard of comparison* lies *in the middle* of the gap. Because of this, the gap is fully determined by two contextual parameters, the standard of comparison, which settles that position of the gap, and the tolerance level, which settles its size: where i is the value of the standard of comparison and k is the tolerance level, the gap is simply $(i + k/2, i - k/2)$. Conversely, by these stipulations, given the upper and lower end points of the gap, both tolerance level and standard of comparison can be recovered.

We assume a basic domain D of individuals, and a domain C of contexts. We shall let each atomic n-place predicate, for $n \geq 1$, be associated with an $n+1$-place function from $D^n \times C$ to the set of truth values $\{0,1\}$. However, the vagueness of many-place predicates has not been much discussed in the literature, and I too shall restrict the discussion to one-place predicates, among the atomic ones.

We assume a language L with a finite stock L_A of atomic one-place predicates. We also assume a classical language with negation, conjunction, and universal quantification as primitive, and the others classically defined as usual. We also have a set L^* of *fragments* of L, where each fragment $l \in L^*$ has the same basic vocabulary as L except for containing only a subset of the set of closed singular terms of L and a subset of the set of vague one-place predicates of L. Each fragment is closed under the same syntactic operations as L itself. We further have a *fragment function* $\mathscr{F} : C \longrightarrow L^*$ from contexts to fragments of L. Let $\mathscr{F}_0(c)$ be the set of one-place predicates in $\mathscr{F}(c)$ and s^0 be the set of one-place predicates in the sentence s. Then we shall require that

(FP) If s is used in context c, then $s^0 \subseteq \mathscr{F}_0(c)$

With each atomic predicate $F_i \in L_A$ is associated a weak order $\lesssim_{F_i}$ and an admissible (at least interval scale) measure function $\mathscr{H}_{F_i}$. We further have a function $\mathscr{T}_c(F)$ that for a context c and an atomic predicate F returns a *tolerance level* for F at c. We also have a *central gap function* $\mathscr{G}_c(F)$ from a predicate F and a context c to pairs (i, j) of real numbers, $i \geq j$, with respect to the measure function $\mathscr{H}_F$. If F is vague, then $i > j$. We shall refer to the left member as $\mathscr{G}_c^+(F)$ and to the right member as $\mathscr{G}_c^-(F)$.

To maintain consistency we need the condition that:

(GT) For every predicate $F \in L$ and context $c \in C$ it holds that
$\mathscr{G}_c^+(F) - \mathscr{G}_c^-(F) \geq \mathscr{T}_c(F)$

That is, the central gap of the context must be at least as great as the tolerance level of the context. As already suggested, however, it will be convenient to let the size of the central gap *equal* the tolerance level, since then an increase of precision, i.e. a reduction of the tolerance level, automatically reduces the central gap. Hence, we stipulate that:

(GT$^=$) For every predicate $F \in L$ and context $c \in C$ it holds that

$$\mathscr{G}_c^+(F) - \mathscr{G}_c^-(F) = \mathscr{T}_c(F).$$

When several atomic predicates are used in the same context $c \in C$, it will be required that individual terms do not have referents with measures in *any* of the gaps associated with the predicates used in c, hence that they have referents outside $\nabla(c, \overline{F})$.

Let f, f' etc. be assignment functions from variables $x_1, x_2, \ldots$ in L to objects in D. The semantic function $[\![\cdot]\!]$ maps a sentence s, a context c and an assignment function f onto a truth value in $\{0,1\}$. Let $[\![F]\!]_c^+ = \{a : \mathscr{H}_F(a) > \mathscr{G}_c^+(F)\}$, and let $[\![F]\!]_c^- = \{a : \mathscr{H}_F(a) < \mathscr{G}_c^-(F)\}$. For a singular term t, let $[\![t]\!]_f$ be $[\![t]\!]$ if t is a constant, and $f(t)$ if t is a variable. The relation $f'[x]f$ holds iff the assignment f' differs from f at most in what it assigns to x. Then we have the truth definition:

(GS) i) Where t is a constant, $[\![t]\!] \notin \nabla(c, \overline{F})$

ii) Where F is an atomic one-place predicate,

$[\![Ft]\!]_{f,c} = 1$ iff $[\![t]\!]_f \in [\![F]\!]_c^+$

iii) $[\![\neg A]\!]_{f,c} = 1$ iff $[\![A]\!]_{f,c} = 0$

iv) $[\![A \& B]\!]_{f,c} = 1$ iff $[\![A]\!]_{f,c} = 1$ and $[\![B]\!]_{f,c} = 1$

v) $[\![Qx(Ax, Bx)]\!]_{f,c} = 1$ iff

$[\![Q]\!] f'(f'[x]f$ and $[\![Ax]\!]_{f',c} = 1$ and $f'(x) \notin \nabla(c, \overline{F}), [\![Bx]\!]_{f',c} = 1)$

According to (GSv), in a context c a sentence $Q(Ax, Bx)$ is true just in case $[\![Q]\!]$ many individuals of which Ax is true in c and that do not belong in the contextually determined cut $\nabla(c, \overline{F})$ are such that Bx is true of them as well.

We can verify that a tolerance principle such as

(13) $\text{Some}x(\text{man}(x) \,\&\, \text{height}(x) \leq n + k \text{ mm}, \text{tall}(x)) \quad \rightarrow$

$\text{All}x(\text{man}(x) \,\&\, \text{height}(x) \geq n \text{ mm}), \text{tall}(x)$

is true in any context *c* where the tolerance level for 'tall' is *k* mm or greater. We assume that the relevant measure function maps men on their heights in mm. Let us assume here that 'man' is non-vague. Then the antecedent of (13) is true just in case some individual *a* in the restricted domain is a man and has a height above the upper edge of the 'tall' gap in *c*. Then any individual *b* in the *c*-restricted domain that is a man and has a height at least that of the height of *a* minus *k* mm, itself has a height above the upper edge of the gap, for there is no individual in the restricted domain that has a height in the gap, and any individual *c* in the domain with a height below the gap is more than *k* mm shorter than *a*. So the the consequent of (13) is true.[14]

7. Some consequences

In this section I shall consider five consequences of the gap semantics that *prima facie* are negative for the account.

7.1. Filling the gap

The function $\mathscr{G}$ is designed to determine a central gap in a context *c* so that no referent of a term used in *c* belongs to the domain cut that results. Moreover, $\mathscr{G}$ is to satisfy *charity* as far as is reasonable ((GAPiv)). And everything explicitly

[14] As mentioned above, similar accounts have been proposed by Ruth Manor, Haim Gaifman, Robert van Rooij, and Mario Gómez-Torrente. On Manor's account (2006), sentences with vague predicates are used with the *presupposition* that there is a gap that is large enough. In contexts where the presupposition is true, the standard of comparison is determined to be in the gap, so that objects both in the extension and in the anti-extension have measures outside. In contexts where the presupposition is false, sentences involving the predicate are all considered false (because of the existence of a gap is stated with large scope in the logical form). Gómez-Torrente's account (2010) is similar to Manor's but more radical: on occasions of use where speaker preconceptions of tolerance, clear cases and domain of individuals does not induce a suitable gap in the domain, the relevant vague predicate lacks extension and intension, and so the sentence uttered lacks truth-conditions.

Gaifman's account (2002) is similar to these in that it relies on the existence of actual gaps. It differs in that in contexts where there *isn't* any gap for a particular tolerant predicate, the predicate simply loses its tolerance and thus becomes sharp. van Rooij (this volume) differs from Manor and Gaifman, in that the gap that is required belongs in a contextually relevant *comparison class*.

These alternatives have serious problems. In Manor's case, the drawback is that all clear case applications of tolerant predicates counterintuitively come out false in case there is no actual gap. In Gómez-Torrente's case, they lack truth-conditions. In Gaifman's case, although he claims that tolerance is a consequence of meaning-determining rules (p. 18), it simply disappears in case non-linguistic circumstances aren't accommodating, a combination of views that I find somewhat bizarre. van Rooij's account does not have such obvious problems, and works well for singular predications. It does not, however, in general work for quantification, since if the quantifier domain itself contains a sorites sequence, the comparison class does not help.

talked about, e.g. by direct enumeration of individuals, must count. As a result, so many objects might be mentioned that no gap that agrees with the verdicts of the speaker still exists. Then the domain may contain a sorites sequence of individuals $d_1, \ldots, d_n$ such that d_0 has a height above the gap initially given as the discourse starts, d_n has a height below the initial gap, and between each d_i, d_{i+1} the difference in height is at most that of the tolerance level.

This is then a sorites sequence with respect to the current tolerance level τ and one particular standard of comparison κ that best agrees with the height verdicts of the speaker. The result will be a sorites contradiction. As $\mathscr{G}$ is specified, a standard of comparison must be chosen that makes one or more of the speaker's verdicts come out false, but preserves consistency, such that the gap is placed above the upper end of the measures of the sorites sequence, or below the lower end. The speaker will then be represented as mistaken in her verdicts about many of the individuals in the domain.

An alternative scenario makes use of the requirement of respecting conservative context updates. This may happen e.g. by various anaphoric constructions that enforce a conservative context update. One possibility is verb phrase ellipsis, as considered by Jason Stanley (2003) in his objection to contextualism about vagueness. We can have a discourse by a speaker S that evolves in the following manner:

(14) John is tall. And so is Alice, and Noah, and ..., but not ..., although ... etc.

which intuitively demands that all the individuals named belong in the domain talked about. Clearly it can happen that the resulting domain contains a sorites sequence. The result is again a contradiction. In this case, however, the value of the $\mathscr{G}$ function is pragmatically restrained by the discourse, and must be the same at the end as in the beginning. But that value is not admissible. As a result, no admissible interpretation can be given. The speaker must be represented as not fullly interpretable.

In some cases there will be a trade-off between representing the speaker as guilty of errors of estimates, or of inconsistency. This seems to me to be perfectly in order, given that the contextual tolerance level really reflects the disposition of the speaker to accept tolerance principles, and thereby reflects the speaker's lack of precision.

7.2. Emptying the domain

An opposite negative effect of the semantics, together with lack of precision of the speakers, is that if many vague predicates are used in some particular context c, and the lack of precision of the relevant speakers enforces a considerable domain cut for each or many of those predicates, then the result may be a virtually empty contextual domain, contrary to what the speakers in the context may think about it.

Maybe this result is intuitively acceptable to a somewhat lower degree. On the other hand, this consequence, as well as the consequence above, would be expected to be rarely exemplified, since it relies on a combination of very low precision, the use of many vague predicates, and pragmatic discourse features that forces these predicates to be treated as used in one and the same context.

7.3. Overly exact statements

Suppose we have an initially given domain, such as the domain of Swedish male citizens between 20 and 60 years old. Suppose this domain contains exactly 2.500.000 individuals. Suppose speaker Nils in context c says

(15) Exactly 1.600.000 Swedish male citizens between 20 and 60 years old are tall.

Suppose further that Nils speaks with a tolerance level of 20 mm with respect to 'tall' in c. Now, we can be sure that such a large domain as this will be tight with respect to that tolerance level, and hence contain a sorites sequence. Because of this, the central gap determined by the gap function $\mathcal{G}$ and the contextual parameters will induce a non-empty cut in the domain. Say that the cut $\nabla(c, \mathit{tall})$ contains 50.000 individuals.

Now, the salient reading of (15) is that which entails

(16) For any number k, if k is the number of Swedish male citizens between 20

and 60 years old, then exactly $k - 1.600.000$ Swedish male citizens between 20 and 60 years old are not tall.

The result is that the sum of those that are said to be tall and those that are entailed by this not to be tall, is larger than the size of the *restricted* domain, which has a size of 2.450.000 individuals.[15] Hence, under the domain restriction interpretation, (15) is false. Either there are fewer tall individuals in the restricted domain than what is explicitly claimed, or there are fewer non-tall individual in the restricted domain than what is entailed by (15) together with background truths, or both.

I think this an intuitively correct consequence. Exact claims about such large populations should not be made with such a high level of tolerance. With high tolerance levels you can make statements that are capable of being approximately true, but claims of such exact character will turn out false.

Sentences like (15) will in general pose problems for contextualist accounts of vagueness, such as Raffman 1994 or 1996, or Fara 2000. They will be true only if there is an overly sharp boundary between two individuals of almost the same height, and those two individuals are identified by the sentence, e.g. number 1.900.000 and number 1.900.001 in the height ranking of Swedish male citizens between 20 and 60 years old. Such judgments are not deemed acceptable by contextualist accounts, and so there must be mechanisms that move the boundary elsewhere, and hence the claim must be considered false. It cannot be the ordinary mechanism (by which considering two highly similar individuals induces a boundary shift so that the two individuals are placed in the same category), simply because no two highly similar individuals are considered. Moreover, the shifting usually appealed to is supposed to be *true*-making, not *false*-making. In some contextualist accounts, such as Shapiro's, the speaker who asserts (15) must be considered *incompetent*.

But it has a peculiar status also on other accounts. On an epistemicist account, (15) is unknowable (even if true), unless one employs a very precise measuring

[15] If you dispute that (16) is a consequence of (15), then complicate (15) with (16) as a second conjunct.

method. On a supervaluationist account, (15) is necessarily not super-true (since that is inconsistent with the existence of truth value gaps), but is still not super-false, since it is true under one sharpening (assuming the example is in the borderline area). On a degree theoretic account, it must have a truth value close to that of a contradiction, since it entails a conjunction $p \& \neg q$ concerning the tallness status of individuals number 1.900.000 and number 1.900.001 in the height ranking, where p and q must have almost the same truth value.

7.4. Complication of inference

A feature of the gap account is the need for tracking the use of vague predicates across contexts. This complicates the account of overt reasoning. Consider a discourse with the following three sentences uttered in sequence:

(17) a. Most men are rich
  b. Most men are tall
  c. Hence, some man is rich and tall

Suppose we are considering an initially given domain where even with a reasonable richness cut, (17a) is true. Similarly, even with a reasonable tallness cut (17b) is true. The third sentence, (17c), is a logical consequence of the first two. However, it contains both 'rich' and 'tall', and so the richness cut and the tallness cut must be combined for evaluating (17c). But this may have the effect of reducing the domain so that some rich men go out because of the tallness restriction and some tall men go out because of the richness restriction. It may then happen that this leaves the domain without any man that is both rich and tall. Hence, in the context, (17c) is false.

This shows that to get the logic right, we must somehow ensure that the same domain is considered across the discourse. This is troublesome in cases where the required restriction is effected *retroactively*: a sentence occurring later in the discourse restricts the domain for a sentence that has occurred earlier, for the sake of the inference.

Note, however, that it happens only to the extent that the earlier sentence is used as a premise in reasoning that is performed at the later stage. It is standard in reasoning that one must check that premises used at an earlier stage of a discourse employ words without equivocation compared with uses in later premises or the conclusion, that context dependent expressions are used with the same values across the discourse etc. Therefore, when using an earlier stated premise, one needs to check that it still holds in the later context, and perhaps reaffirm in it in the new context. From this perspective, the complication is not radically new.

7.5. Content preservation

The last consequence I shall consider here is the one I see as potentially most worrying. It concerns whether the intuitive *utterance content* is preserved by the gap account.

Normally, when domain restriction is introduced in natural language semantics, it is introduced to get the intuitively perceived content right. For instance, in

(11) Everyone left at midnight

we introduce the domain restriction to arrive at the intuitive content that everyone *at the party* left at midnight.

In the gap semantics, it is not so clear that the domain restriction tracks the intuitive content of the utterance. When the average speaker Nils says

(18) Most Swedish women are tall

it is not likely that he *intuitively* wants to convey that most Swedish women in a *suitably restricted* subset of the set of Swedish women are tall. And if not, it appears that the semantics simply gets the intended content wrong.

The intuition behind this objection is pretty strong. However, it also seems to that, on reflection, there is more to be said on this matter in defense of the proposal. First, it is rather clear, I think, that the speaker Nils in this case does speak with non-maximal precision. He does not have a sharp demarcation of tall women, or tall Swedish women, in mind. It would therefore *not* be fully accurate to represent Nils as saying something equivalent to

(19) For any number k, at the current time t, if k is the number of Swedish women at t, then the number of *tall* Swedish women at t is greater than $k/2$.

Rather, it seems intuitively right to say that Nils has some kind of *typicality* judgment in mind: it is *typical* of Swedish women that most of them are tall. Since it is typical but not precise, it should be *approximately* true. And taking it as approximately true is taking it as literally true *with respect to* an approximation. Again, it seems reasonable to equate the idea of an approximation with the idea of a standard of precision, and therefore in the case of vagueness, with the tolerance level. Statement (18) is meant to be true with respect to the accepted imprecision of 'tall'.

The next step is to ask how *truth* is approximated by means of a tolerance level. *One* way of approximating truth by way of the tolerance level is to make a statement that is *literally* true with respect to a domain that is approximately the same is the original domain, and where the difference is induced by the standard of precision.

On this way of seeing the matter, the gap semantics offers one way of spelling out the approximation idea such that the intuitive truth of tolerance principles is preserved.

Department of Philosophy

Stockholm University

References

Austin, John Langshaw (1975). *How to do Things with Words*. 2nd ed. Oxford: Oxford University Press.

Fara, Delia Graff (2000). 'Shifting sands: an interest-relative theory of vagueness'. In: *Philosophical Topics* 28, pp. 45–81.

Gaifman, Haim (2002). 'Vagueness, Tolerance, and Contextual Logic'. Unpublished Ms.

Gómez-Torrente, Mario (2010). 'The Sorites, Linguistic Preconceptions, and the Dual Picture of Vagueness'. In: *Cuts and Clouds*. Ed. by Richard Dietz and Sebastiano Moruzzi. Oxford: Oxford University Press, pp. 228–53.

Keefe, Rosanna and Peter Smith (1996). 'Introduction: theories of vagueness'. In: *Vagueness: A Reader*. Ed. by Rosanna Keefe and Peter Smith. MIT Press, pp. 1–57.

Kennedy, Christopher (1999). 'Gradable Adjectives Denote Measure Functions, not Partial Functions'. In: *Studies in the Linguistic Sciences* 29.1.

— (2007). 'Vagueness and Grammar: the Semantics of Relative and Absolute Gradable Adjectives'. In: *Linguistics and Philosophy* 30, pp. 1–45.

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

Krantz, David H et al. (1971). *Foundations of Measurement*. Vol. I. Republished by Dover Publications, NY, 2007. Academic Press.

Lewis, David (1979). 'Scorekeeping in a Language Game'. In: *Journal of Philosophical Logic* 8. Reprinted in Lewis 1983, 233-49. Page references to the reprint, pp. 339–59.

— (1983). *Philosophical Papers. Volume I*. Oxford: Oxford University Press.

Manor, Ruth (2006). 'Solving the Heap'. In: *Synthese* 153, pp. 171–86.

Pagin, Peter (2005). 'Compositionality and context'. In: *Contextualism in Philosophy*. Ed. by Gerhard Preyer and Georg Peter. Oxford: Oxford University Press, pp. 303–48.

— (2010). 'Vagueness and Central Gaps'. In: *Cuts and Clouds*. Ed. by Richard Dietz and Sebastiano Moruzzi. Oxford: Oxford University Press, pp. 254–72.

Raffman, Diana (1994). 'Vagueness without Paradox'. In: *Philosophical Review* 103.1, pp. 41–74.

— (1996). 'Vagueness and Context-Relativity'. In: *Philosophical Studies* 81, pp. 175–92.

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

Stanley, Jason (2003). 'Context, interest relativity and the sorites'. In: *Analysis* 63, pp. 269–80.

Stanley, Jason and Zoltan Szabó (2000). ‘On quantifier domain restriction’. In: *Mind and Language* 15, pp. 219–61.

Westerståhl, Dag (1985). ‘Determiners and Context Sets’. In: *Generalized Quantifiers in Natural Language*. Ed. by Johan van Benthem and Alice ter Meulen. Amsterdam: Foris, pp. 45–71.

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

Wright, Crispin (1976). ‘Language-mastery and the sorites paradox’. In: *Truth and Meaning*. Ed. by Gareth Evans and John McDowell. Oxford: Clarendon Press, pp. 223–47.

Zardini, Elia (2008). ‘A Model of Tolerance’. In: *Studia Logica* 90, pp. 337–68.