Vagueness and Central Gaps*

Peter Pagin

October 13, 2008

Abstract

Ordinary intuitions that vague predicates are *tolerant*, or cannot have sharp boundaries, can be formalized in first-order logic in at least two non-equivalent ways, a stronger and a weaker. The stronger turns out to be false in domains that have a *significant central gap* for the predicate in question, i.e. where a sufficiently large middle segment of the ordering relation (such as *taller* for 'tall') is uninstantiated. The weaker principle is true in such domains, but does not in those domains induce the sorites conclusion.

This fact can be used for interpreting ordinary uses of vague expressions by means of a new kind of contextual quantifier domain restriction. A central segment is cut from the domain, if consistent with speaker intentions. As long as this is possible, tolerance, bivalence and consistency can all be retained.

This paper focuses on the basic semantic properties in a model-theoretic setting. The natural language application is sketched and the nature of the approach briefly discussed.

1 Tolerance principles

A sorites argument in the inductive format is normally taken to have the following form:

(1)
1 $F(k_1)$
2 $\forall i(F(k_i) \rightarrow F(k_{i+1}))$
―――――――――
3 $F(k_n)$

*This paper was presented at the 6th Arché vagueness workshop in St Andrews, in March 2006, at the Logic and Language conference in Birmingham, in April 2006, and at the Stockholm Logic and Language seminar in September 2006. I am grateful to the audiences on those occasions, in particular to Herman Cappelen, Manuel Garcia-Carpintero, Richard Dietz, Patrick Greenough, Patrick Grim, Jeff King, Andrew McGonigal, Augustin Rayo, Tim Williamson, and Crispin Wright. I owe especially much to Sven Rosenkranz who was my commentator in Birmingham, and who also provided me with further helpful comments on a later version. I also owe much to Kathrin Glüer for discussions of these ideas over several years. The work on this paper was funded by a research grant from The Swedish Research Council.

The terms '$k_1, \ldots k_n$' are taken to denote objects $a_1, \ldots a_n$ in a sequence A along some ordering relation, such as *taller than*. We have an instance of the sorites *paradox* if the premises are apparently true while the conclusion is apparently false. We get a typical example by choosing 'is tall for a man' for 'F', and $a_1, \ldots a_n$ as a sequence of men so that a_1 is 200 cm tall, and each a_{i+1} is 1 mm shorter than his predecessor, while a_n is 150 cm tall. In such a case both premises do seem true, while the conclusion also does seem false.

The argument is inductive, and the second premise is the *inductive premise*. The apparent acceptability of the inductive premise derives from a basic intuition about typically vague terms: that they do not have sharp boundaries. In the case imagined, that intuition would typically be expressed by means of

(2) 1 mm cannot make the difference between being tall and not being tall.

Put in a more inductive format, we would say, in this case:

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

When put in the format of 3, the formulation is apt for expressing a version of the basic intuition: that the predicate in question is *insensitive to small differences*. The property of being insensitive to small differences is what Crispin Wright has called 'tolerance' (Wright 1976, 156). 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.

For reasons that are briefly given later (section 5), I prefer to speak of linguistic expressions as being vague or tolerant, rather than concepts, but that is for now less significant.[1] Wright's introduction of the notion of tolerance, over and above the notion of lacking sharp boundaries, was important, not just because a term may *in fact* lack sharp boundaries for one reason or other,[2] but because the tolerance of a vague term intuitively *explains* why it lacks sharp boundaries. It is conceivable that a term lack sharp boundaries in all worlds where it has the same meaning, even though there is a nomological or metaphysical explanation of why this is so that has nothing to do with vagueness.[3]

[1] I also find they claim, that the justice with which the term applies isn't affected, unnecessarily strong. It would be enough for purposes of an account of vagueness, it seems to me, if the degree of justification drops by approximating a limit value that is high enough.

[2] This could be, for instance, because by nomic regularity, animal species cannot be very similar. As long as we treat vagueness extensionally, it can also be because the predicate has an empty extension.

[3] We cannot explain in general how the distinction between being true in virtue of meaning properties and being true in virtue of properties of what is denoted is to be applied (there

For present purposes I shall refer to principles like 3 as *tolerance* principles, even though, for the reasons just stated, they can be true because of other factors than insensitivity. A tolerance principle is an inductive principle that involves a tolerant predicate —'tall' in the case of 3 — and a *tolerance level* — 1 mm in the case of 3. A tolerance level, as I use the term, then, does not depend on any particular sequence of objects, but makes explicit what the difference is to be in a sequence of objects that is appropriate for an inductive premise like the one in 1.

Since tolerance principles are themselves inductive principles we can as well state a sorites argument with their help, without relying on any particular sequence. In order to so, we must make the logical form precise. It quickly turns out, however, that the informal statement 3 admits of two different formalizations. The form of antecedent and consequent,

(4) A man of x mm is tall

as it occurs in 3, has a *generic* interpretation, which, I assume, is equivalent to

(5) $\forall y(\text{Man}(y) \ \&\ H(y) \geq x \ \rightarrow \text{Tall}(y))$

with 'H' for 'the height of ... in mm'.[4] Accordingly, given that 3 itself is of conditional form, the natural predicate logic rendering would be

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

On the other hand, we would also take 3 to have a reading like this:

(6) For any two men, if the one is (not more than) $n+1$ mm tall and the other is (at least) n mm tall, then if the former is tall, so is the latter.

And, the natural formalization of 6 is

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

But although we seem to get (T1) as well as (T2) out of the intuitive formulation 3, (T1) and (T2) are not equivalent. The difference is made more perspicuous by noting that (T2) is equivalent with

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

---

does not seem to be an analytic/synthetic dichotomy in this sense), but we may be able to do somewhat better when it comes to tolerance.

[4]The generic reading of 4 involves an element of *nomicity* that is not captured by the quantification in 5, but I shall not concern myself with that aspect.

I shall call tolerance principles of the form (T1) *strong tolerance principles*, and those of the form (T2)/(T2′) *weak tolerance principles*.[5] Weak and strong tolerance principles work differently. With respect to some cases they induce the same result, but not in all. Consider the original setup, with the sequence $A = \langle a_1, \ldots, a_n \rangle$, and

$$\text{Let } G(n) =_{def} \forall y(H(y) \geq n \ \rightarrow \text{Tall}(y))$$

We then get with (T1) a modified sorites argument:

(7)	1	$G(2000)$	assump.
	2	$H(k_n) = 1500 \ \& \ \neg \ \text{Tall}(k_n)$	assump.
	3	$\forall i(G(i+1) \ \rightarrow \ G(i))$	(T1)
	4	$G(1500)$	1, 3
	5	$\text{Tall}(k_n)$	2, 4
	6	$\perp$	2, 5

Note that the absurdity conclusion requires the assumption 2, for by itself, 4, i.e.

$$\forall y(H(y) = 1500 \ \rightarrow \text{Tall}(y))$$

is true if there are no men of 150 cm height in the universe of discourse.

Similarly, 1 is true if there are no men of height 200 cm in the universe of discourse. We could add a dependence on actual instantiation and drop the assumption 1. We then replace it with assumption about k_1, together with a *number uniformity* principle, stating that anyone is tall whose height is at least equal to that of someone who is tall[6]:

(U) $\forall x, y(H(y) \geq H(x) \ \rightarrow \ (Tall(x) \rightarrow Tall(y)))$

We then have the simple derivation

(8)	1	$H(k_1) = 2000 \ \& \ \text{Tall}(k_1)$	assump.
	2	$G(2000)$	1, (U)

which can replace the first assumption in 7. In that case the contradiction depends on the (U), (T1), and the facts about the heights and tallness attributes of k_1 and k_n. In either case, at least one singular fact is needed.

If instead we use the weaker tolerance principle (T2), the derivation will be rather different. First, another abbreviation:

[5] I shall henceforth take quantification to be restricted to a domain of men, and so the 'Man' conjunct will be dropped.

[6] This is of course a simplification, since a man of 200 cm height and 250 cm width would not count as tall, but I shall disregard shape. Note that (U) implies both the claim that men of the same height have the same tallness attribute, and the so-called penumbral principle that anyone taller than someone who is tall, is tall. The penumbral principle follows on the assumption that x is taller than y iff $H(x) > H(y)$.

Let $E(n) =_{def} \exists x(H(x) \leq n \ \& \ \text{Tall}(x))$

With respect to the same domain of men, we will then have the sorites argument:

(9)	1	$H(k_1) = 2000 \ \& \ \text{Tall}(k_1)$	assump.
	2	$H(k_n) = 1500 \ \& \ \neg\text{Tall}(k_n)$	assump.
	3	$\forall i(E(i+1) \ \rightarrow \ G(i))$	(T2)
	4	$E(2000)$	1
	5	$G(1999)$	(T2), 4
	6	$H(k_2) = 1999 \ \rightarrow \ \text{Tall}(k_2)$	5
	7	$H(k_2) = 1999$	assump.
	8	$\text{Tall}(k_2)$	6, 7
	9	$H(k_2) = 1999 \ \& \ \text{Tall}(k_2)$	7, 8
	10	$E(1999)$	9
	11	$G(1998)$	10, (T2)
		$\vdots$	
	12	$G(1500)$	
	13	$H(k_n) \geq 1500 \ \rightarrow \ \text{Tall}(k_n)$	12
	14	$\text{Tall}(k_n)$	2, 13
	15	$\neg\text{Tall}(k_n)$	2
	16	$\perp$	14, 15,

For deriving the contradiction with (T2), we need as premises both the heights and tallness attributes of a positive and a negative specimen (i.e. premises 1 and 2 of sorites argument 9), but also each of the height instantiations (including 7). That is, the sorites argument with (T2) requires that there be an unbroken sequence of individuals from a positive instance to a negative instance, where adjacent members differ by at most the tolerance level (in the example, 1 mm).

Initially, this may seem to mean nothing more than that a simple and elegant sorites argument can be made complex and cumbersome. However, the difference between (T1) and (T2) gains importance in relation to a domain where the chain in question has a *significant central gap.* Informally, a significant central gap is a gap in the chain which is bigger than any admissible tolerance level, and such that there are only positive instances on the one side of the gap and only negative instances on the other side. With respect to our domain of men of different heights, we would have such a gap if, say, all members of heights between 190 cm and 160 cm were removed.

Thus, suppose that from the vocabulary, the terms $k_{102}, \ldots, k_{400}$, denoting men shorter than 190 cm and taller than 160 cm, are removed. What happens with the sorites arguments? As regards (T1), the argument 7 goes through as before, since (T1) only requires one negative instance (and in the expanded version, one positive and one negative instance). As regards (T2), however, the argument 9 clearly does *not* go through: in order to derive the intermediate conclusion that all men of a height of 1898 mm are tall, we need as a premise that there is a man of at most 1899 mm who is tall, but since this time there are no men in the domain of heights between 1900 mm and 1898 mm, the derivation is broken at this step. The instantiation picks up again at the term k_{401}, denoting a man of 160 cm height, but that man is already clearly short, and so the premise Tall(k_{401}) can not be justifiably added. Therefore, the derivation cannot be resumed at this step, and because of that the conclusion Tall(k_n) cannot be reached. That is, because the tolerance principle that would be needed to bridge the gap postulates an unacceptably high tolerance level—30 cm—the conclusion cannot be reached, and the sorites argument fails.

That the (T2) argument fails is not of great significance if the (T1) argument, which does not rely on an unbroken chain of instantiations goes through anyway. However, it doesn't. Interpreted in a domain with a significant central gap, the (T1) principle is straightforwardly *false*. With respect to our example, consider the instance of (T1) at the lower end of the gap:

(10) $G(1601) \rightarrow G(1600)$

i.e.

10 $\forall y(H(y) \geq 1601 \rightarrow \text{Tall}(y)) \rightarrow \forall y(H(y) \geq 1600 \rightarrow \text{Tall}(y))$

To make things explicit, we add as an extra premise

(11) $\neg$ Tall(k_{401})

where the term 'k_{401}' denotes the man of 1600 mm height. Now, the antecedent of 10 is true: anyone in the domain that is 1601 mm or higher is 1900 mm or higher, and so tall. The consequence, however, is false: there is an individual of 1600 mm, the referent of 'k_{401}', and that individual is not tall, by the truth of 11. Hence, the instance 10 of (T1) is false, and so (T1) itself is false. Therefore, in a gappy domain, the falsity of (T1) does *not* entail the existence of a sharp boundary.

By contrast, (T2) is true at the lower end of the gap, for

(12) $\exists y(H(y) \leq 1601 \ \& \ \text{Tall}(y)) \rightarrow \forall y(H(y) \geq 1600 \rightarrow \text{Tall}(y))$

is true: the antecedent is false, since any domain member shorter than or equal to 1601 mm is shorter than or equal to 1600 mm, and therefore not tall. But since the antecedent is false, we can't detach the consequent, and don't have a contradiction.

Similarly, (T2) is true at the upper end of the gap.

(13) $\exists y(H(y) \leq 1900 \ \& \ \text{Tall}(y)) \rightarrow \forall y(H(y) \geq 1899 \rightarrow \text{Tall}(y))$

The antecedent is true, but so is the consequent, for any domain member with a height of at least 1899 mm has a height of at least 1900 mm, and so is tall.

So with respect to a domain with a significant central gap, the strong tolerance principle (T1) is false, but the falsity does not entail that there is a sharp boundary. The weak tolerance (T2) is true, but does not lead to inconsistency. Only the existence of a sharp boundary makes (T2) false (see next section). Again, these facts might seem to be of marginal interest, since the difficult problems concerning vagueness relate to non-gappy domains. I shall, however, make use of these facts for a semantics of natural language sentences where the existence of significant central gaps is not presupposed.

In the next section, the (T1) and (T2) principles, and their semantic relation, will be characterized model-theoretically. It will turn out that under natural conditions, (T2) is a consequence of (T1). The non-technical reasoning, and the application to vagueness, will be resumed in section 3.

2 Models of tolerance and gaps

To characterize the (T1) and (T2) principles formally, we define a class of models:

Definition 2.1. A *V–model* $\mathcal{M}$ is a classical model $\langle \mathcal{D}, \mathcal{I}, \mathcal{T}, \mathcal{H} \rangle$ for a first-order language $\mathcal{L}$, where $\mathcal{D}$ is a domain of individuals and $\mathcal{I}$ an interpretation function, $\mathcal{T}$ a function that for each predicate letter F in $\mathcal{L}$ as argument assigns a real number $\mathcal{T}_F \geq 0$, and $\mathcal{H}$ a function that for each predicate letter F in $\mathcal{L}$ assigns a total function $\mathcal{H}_F$ from $\mathcal{D}$ to $\mathcal{R}$, such that

$$(+) \qquad \text{for any } a, b \in \mathcal{D}\,(\mathcal{H}_F(a) \geq \mathcal{H}_F(b) \rightarrow \; (b \in \mathcal{I}(F) \;\rightarrow\; a \in \mathcal{I}(F)))$$

By a *classical model* I mean a model where the truth definition gives the standard classical clauses for the first order logical constants. These clauses will be assumed below. The function $\mathcal{T}_F$ will be used below for assigning a tolerance level to the predicate F in the model.

Definition 2.2. A V-model $\mathcal{M}$ *has a central F-gap* iff there are no objects $a, b \in \mathcal{D}_{\mathcal{M}}$ such that $a \in \mathcal{I}(F)$, $b \notin \mathcal{I}(F)$ and $0 \leq \mathcal{H}_F(a) - \mathcal{H}_F(b) \leq \mathcal{T}_F$

Remark 2.3. It would be natural to identify a central gap with a pair $\langle i, j \rangle$ of real numbers such that i is the least upper bound and j the greatest lower bound of the gap. But although this can be done, there is no simple and uniform definition of the pair. For if we allow dense and continuous domains of objects, there need not be any smallest member of the set (for instance, for every long period of time in the domain, there may be a shorter period of time in the domain that is still long, even though not every period is), nor any largest number that is not the measure of a member (e.g. for every real number u that is not the length of a long period in the domain, there is greater number v that is also not the length of a long period in the domain, even though there are long periods). We are only guaranteed by the Least Upper Bound Axiom of real analysis that one of the two must exist. But then to identify a gap with a pair of numbers leads to a considerable increase in the complexity of definitions and proofs.

Secondly, as the definition is stated, a domain has a central F-gap even in case all objects are F:s or no object is. Again, it simplifies the definition and the reasoning with it not to require both positive and negative instances in the domain. But if we don't, there need not even be a definite pair $\langle i, j\rangle$ to identify with the gap.

However, for the natural language semantics to be presented in section 4, it is more convenient to identify gaps with definite pairs of numbers.

End of remark.

Now we can characterize the two tolerance principles model theoretically:

Definition 2.4. A (T1)-model $\mathcal{M}$ for a predicate F of $\mathcal{L}$ is a V-model $\langle \mathcal{D}, \mathcal{I}, \mathcal{T}, \mathcal{H}\rangle$ for $\mathcal{L}$ such that $\mathcal{T}_F > 0$, and

(ti) For all k (if (for all $a \in \mathcal{D}_\mathcal{M}$ (if $\mathcal{H}_F(a) \geq k + \mathcal{T}_F,$ then $a \in \mathcal{I}(F)$), then (for all $a \in \mathcal{D}_\mathcal{M}$ (if $\mathcal{H}_F(a) \geq k,$ then $a \in \mathcal{I}(F)$)))

Definition 2.5. A (T2)-model $\mathcal{M}$ for a predicate F of $\mathcal{L}$ is a V-model $\langle \mathcal{D}, \mathcal{I}, \mathcal{T}, \mathcal{H}\rangle$ for $\mathcal{L}$ such that $\mathcal{T}_F > 0$, and

(tii) For all k, (if (there is $a \in \mathcal{D}_\mathcal{M}$ ($\mathcal{H}_F(a) \leq k + \mathcal{T}_F,$ and $a \in \mathcal{I}(F)$), then for all $a \in \mathcal{D}_\mathcal{M}$ (if $\mathcal{H}_F(a) \geq k,$ then $a \in \mathcal{I}(F)$))

We shall now prove some elementary properties of the (T1) and (T2) models.

Definition 2.6. A V-model $\mathcal{M}$ is

a) $\mathcal{H}_F$-*unbounded* iff for all $k \in \mathcal{R}$ there is $a \in \mathcal{D}_\mathcal{M}$ such that $\mathcal{H}_F(a) \geq k$

b) *empty* iff $\mathcal{D}_\mathcal{M} = \varnothing$

c) F-*full* iff $\mathcal{D}_\mathcal{M} \cap \mathcal{I}(F) = \mathcal{D}_\mathcal{M}$

d) F-*empty* iff $\mathcal{D}_\mathcal{M} \cap \mathcal{I}(F) = \varnothing$

e) F-*free* iff $\mathcal{M}$ has a central F-gap

Fact 2.7. *For all V-models $\mathcal{M}$, $\mathcal{M}$ is a (T1)-model for a predicate F of $\mathcal{L}$ iff $\mathcal{M}$ is empty, or $\mathcal{M}$ is non-empty and F-full, or $\mathcal{M}$ is non-empty and F-empty and $\mathcal{H}_F$-unbounded.*

Proof. Left to right. Assume for reductio that $\mathcal{M}$ is a) non-empty, and b) not F-full, and c) not both F-empty and $\mathcal{H}_F$-unbounded. Because of a) and b) there is a non-F, i.e. an object $b^- \in \mathcal{D}_\mathcal{M} - \mathcal{I}(F)$. Because of a) and c), either – case A – there is an F, i.e. an object $b^+ \in \mathcal{D}_\mathcal{M} \cap \mathcal{I}(F)$, or – case B – $\mathcal{M}$ is $\mathcal{H}_F$-bounded, i.e. there a $k \in R$ such that for all $a \in \mathcal{D}_\mathcal{M}, \mathcal{H}_F(a) < k$.

Consider first case A. By clause (+) of the definition of a V-model, it holds that

(i) for all $a \in \mathcal{D}_\mathcal{M}$ (if $\mathcal{H}_F(a) \geq \mathcal{H}_F(b^+),$ then $a \in \mathcal{I}(F)$)

Then by repeated applications of (ti), from (i) we finally get

(ii) $$\text{for all } a \in \mathcal{D}_{\mathcal{M}} \, (\text{if } \mathcal{H}_F(a) \geq \mathcal{H}_F(b^-), \text{ then } a \in \mathcal{I}(F))$$

Since $\mathcal{H}_F(b^-) \geq \mathcal{H}_F(b^-)$, by (ii) we can conclude that $b^- \in \mathcal{I}(F)$, contradicting the assumption. Hence, case A cannot hold.

Consider then case B. Let $k_j > \mathcal{H}_F(a)$ for all $a \in \mathcal{D}_{\mathcal{M}}$. Then

(iii) $$\text{for all } a \in \mathcal{D}_{\mathcal{M}} \, (\text{if } \mathcal{H}_F(a) \geq k_j, \text{ then } a \in \mathcal{I}(F))$$

Then we can again apply (ti) from (iii) repeatedly, until we have derived (ii), again concluding that $b^- \in \mathcal{I}(F)$, contradicting the assumption. Hence, case B cannot hold either.

Right to left. Three cases: A) $\mathcal{M}$ is empty, B) $\mathcal{M}$ is non-empty and F-full, C) $\mathcal{M}$ is non-empty and F-empty and $\mathcal{H}_F$-unbounded.

Consider case A. Since by assumption $\mathcal{D}_{\mathcal{M}} = \varnothing$ it holds for any $k \in \mathcal{R}$ that

(iv) $$\text{for all } a \in \mathcal{D}_{\mathcal{M}} \, (\text{if } \mathcal{H}_F(a) \geq k, \text{ then } a \in \mathcal{I}(F))$$

Therefore, it also holds for any $k \in \mathcal{R}$ that

(v) $$\text{if for all } a \in \mathcal{D}_{\mathcal{M}} \, (\text{if } \mathcal{H}_F(a) \geq k_j + \mathcal{T}_F, \text{ then } a \in \mathcal{I}(F)),$$
$$\text{then for all } a \in \mathcal{D}_{\mathcal{M}} \, (\text{if } \mathcal{H}_F(a) \geq k_j, \text{ then } a \in \mathcal{I}(F))$$

and then (ti) holds as well.

Case B. Since by assumption it holds for all $a \in \mathcal{D}_{\mathcal{M}}$ that $a \in \mathcal{I}(F)$, (iv) will again hold for any $k \in \mathcal{R}$, and so we have the same conclusion as in case A.

Case C. Assume for reductio that (ti) is false, and hence there is a k_j such that

(vi) $$\text{for all } a \in \mathcal{D}_{\mathcal{M}} \, (\text{if } \mathcal{H}_F(a) \geq k_j + \mathcal{T}_F, \text{ then } a \in \mathcal{I}(F))$$

is true, while

(vii) $$\text{for all } a \in \mathcal{D}_{\mathcal{M}} \, (\text{if } \mathcal{H}_F(a) \geq k_j, \text{ then } a \in \mathcal{I}(F))$$

is false.

Since by the first C-case assumption it holds for any $a \in \mathcal{D}_{\mathcal{M}}$ that $a \notin \mathcal{I}(F)$, (vi) is true only if it holds that

(viii) $$\text{for all } a \in \mathcal{D}_{\mathcal{M}}(\mathcal{H}_F(a) < k_j + \mathcal{T}_F)$$

But since by the second C-case assumption $\mathcal{M}$ is $\mathcal{H}_F$-unbounded, this does not hold. Hence (ti) is true.

In all three cases, (ti) true, and therefore $\mathcal{M}$ is a (T1) model. □

Fact 2.8. *For all V-models $\mathcal{M}$, $\mathcal{M}$ is a (T2)-model for a for a predicate F of $\mathcal{L}$ iff $\mathcal{M}$ is empty, or $\mathcal{M}$ is non-empty and F-full, or $\mathcal{M}$ is non-empty and F-empty, or $\mathcal{M}$ is non-empty and F-free.*

Proof. Left to right. Assume for reductio that a) $\mathcal{M}$ is non-empty, b) not F-full, c) not F-empty, and d) not F-free. Because of a) and b), there is a non-F, i.e. an object $b^- \in \mathcal{D}_\mathcal{M} - \mathcal{I}(F)$. Because of c), there is an F, i.e. an object $b^+ \in \mathcal{D}_\mathcal{M} \cap \mathcal{I}(F)$. Since $\mathcal{M}$ is not F-free, there are two objects $a, b \in \mathcal{D}_\mathcal{M}$ such that $a \in \mathcal{I}(F)$, $b \notin \mathcal{I}(F)$ and $0 \leq \mathcal{H}_F(a) - \mathcal{H}_F(b) \leq \mathcal{T}_F$. Hence

$$\text{(i)} \qquad \mathcal{H}_F(a) \leq \mathcal{H}_F(b) + \mathcal{T}_F$$

Because $a \in \mathcal{I}(F)$, we have

$$\text{(ii)} \qquad \text{there is } c \in \mathcal{D}_\mathcal{M}\, (\mathcal{H}_F(c) \leq \mathcal{H}_F(b) + \mathcal{T}_F, \text{ and } c \in \mathcal{I}(F))$$

Applying (tii) we can conclude from (ii)

$$\text{(iii)} \qquad \text{for all } c \in \mathcal{D}_\mathcal{M}\, (\text{if } \mathcal{H}_F(c) \geq \mathcal{H}_F(b), \text{ then } c \in \mathcal{I}(F)))$$

Instantiating, we have the conclusion that $b \in \mathcal{I}(F)$, contrary to assumption. Therefore, $\mathcal{M}$ cannot be simultaneously non-empty, not F-full, not F-empty, not F-free.

Right to left. We have four cases: A) $\mathcal{M}$ is empty, B) $\mathcal{M}$ is non-empty and F-full, C) $\mathcal{M}$ is non-empty and F-empty, and D) $\mathcal{M}$ is non-empty and F-free.

Case A. By assumption, $\mathcal{M}$ is empty, and so

$$\text{(iv)} \qquad \text{there is } a \in \mathcal{D}_\mathcal{M}\, (\mathcal{H}_F(a) \leq k + \mathcal{T}_F, \text{ and } a \in \mathcal{I}(F))$$

is false for all $k \in \mathcal{R}$, and hence the conditional

$$\text{(v)} \quad \text{if (there is } a \in \mathcal{D}_\mathcal{M}\, (\mathcal{H}_F(a) \leq k + \mathcal{T}_F, \text{ and } a \in \mathcal{I}(F)), \\ \text{then (for all } a \in \mathcal{D}_\mathcal{M}\, (\text{if } \mathcal{H}_F(a) \geq k, \text{ then } a \in \mathcal{I}(F))$$

is true for all $k \in \mathcal{R}$, and so (tii) follows.

Case B. By assumption, $\mathcal{M}$ is non-empty and F-full. Then, all $a \in \mathcal{D}_\mathcal{M}$ are also in $\mathcal{I}(F)$, and so

$$\text{(vi)} \qquad \text{for all } a \in \mathcal{D}_\mathcal{M}\, (\text{if } \mathcal{H}_F(a) \geq k, \text{ then } a \in \mathcal{I}(F))$$

is true for all $k \in \mathcal{R}$. Hence, (v) is again true for all $k \in \mathcal{R}$.

Case C. By assumption, $\mathcal{M}$ is non-empty and F-empty. As in case A, (iv) is false for all $k \in \mathcal{R}$, and the rest follows.

Case D. By assumption, $\mathcal{M}$ is non-empty and F-free. If $\mathcal{M}$ is F-empty or F-full, (tii) follows, so assume $\mathcal{M}$ is neither. In order for (v) to be false for a particular $k \in \mathcal{R}$, there must a pair $\langle b^+, b^- \rangle \in \mathcal{M}$ such that $b^+ \in \mathcal{I}(F)$, $b^- \notin \mathcal{I}(F)$, and

$$\text{(vii)} \qquad \mathcal{H}_F(b^+) \leq \mathcal{H}_F(b^-) + \mathcal{T}_F$$

Then (v) is false for $k = \mathcal{H}_F(b^-)$.

However, if there is such a pair $\langle b^+, b^- \rangle \in \mathcal{M}$, then either $\mathcal{H}_F(b^-) > \mathcal{H}_F(b^+)$, or $0 \leq \mathcal{H}_F(b^+) - \mathcal{H}_F(b^-) \leq \mathcal{T}_F$. The first disjunct is ruled out by condition (+) of the definition of V-models. If the second disjunct is true, then by definition 14, there is no central F-gap in $\mathcal{M}$. Since by assumption there is a central F-gap, the second disjunct is false as well. Then, (v) cannot be false for any k, and hence (tii) holds.

In all four cases, (tii) holds, and therefore $\mathcal{M}$ is a (T2) model. □

Fact 2.9. $(T1) \models_V (T2)$ *but* $(T2) \not\models_V (T1)$

Proof. It follows from Facts 2.7 and 2.8 that the class of (T1)-models is a proper sub-class of the class of (T2)-models. (T1) will be false in (T2)-models $\mathcal{M}$ that are non-empty, not F-full, not F-empty and F-free, as well as in (T2)-models $\mathcal{M}$ that are non-empty, not F-full, F-empty but not $\mathcal{H}_F$-unbounded. □

3 Methodological interlude

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.[7]

I think it is fair to say alternative b) has been the most popular one in the literature of recent decades. Epistemicists such as Timothy Williamson (1994) take the inductive premise to be straightforwardly false, although it is not knowable where the boundary is. Supervaluationists, such as Kit Fine (1975) or Rosanna Keefe (2000), take the inductive premise to be (super-)false, since false in every classical evaluation. Degree theorists, such as Kenton F. Machina (1976), take the inductive premise to be almost completely false, even though each of its instances is almost completely true. Contextualists, like Diana Raffman (1994, 1996), Scott Soames (1999), Delia Graff (2000), and Stewart Shapiro (2003, 2006) take the inductive premise to be false in each context, although the boundary shifts between contexts.

I shall propose that we accept the inductive premise, i.e., in the form of the surviving tolerance principle, exemplified by (T2), or rather, a certain revised version of it. Without going deeply into polemics with the dominant trend, I take it to be intuitively part of the semantics of vague predicates to be insensitive to small differences. Our inability to locate any sharp boundary of vague predicates, and even more the intuitive rejection by ordinary speakers of the very idea of a sharp boundary, suggests that it is part of the meaning

[7] Cf. Keefe and Smith 1996b, 10.

of vague expression as used by ordinary speakers not to have them. Although it is well known that all the theories of vagueness that involve rejecting the inductive premise come at a high cost, rejection in any of the proposed forms may still seem to involve smaller costs than does acceptance.

Since drastic revisions of logic are required to treat the standard sorites argument as invalid, and since it is implausible to deny the existence of tall men, or of non-tall men, the only remaining option seems to be that the use of vague vocabulary is incoherent, or even inconsistent. This position has been advocated e.g. by Michael Dummett (1975). Dummett says

> Wang's paradox merely reflects this inconsistency. What is in error is not the principles of reasoning involved, nor, as on our earlier diagnosis, the induction step. The induction step is correct, according to the rules of use governing vague predicates such as 'small': but these rules are themselves inconsistent, and hence the paradox. Our earlier model for the logic of vague expressions thus becomes useless: there can be no coherent such logic (1975, 265).

A little before that, Dummett provides a gloss on 'consistent':

> 'Consistent' here means that it would be impossible to force someone, by appeal to rules that he acknowledged as correct, to contradict himself over whether the predicate applied to a given object (1975, 264).

That the use of vague vocabulary is inconsistent may be seen as rendered plausible from considerations of so-called *forced march sorites*, which "is designed to force us, one step at a time, into a separate verdict on each successive pair of adjacent items in a sorites sequence" (Horgan 1994, 173). Our inclination to respect initial intuitions about clear cases while not accepting any sharp boundary, naturally leads us into trouble.

On the other hand, there is a frequent use of vague vocabulary that by common sense standards fulfill its communicative function pretty well. In fact, most applications of vague predicates appear *unproblematic*. This observation has been used in an objection against Dummett by Crispin Wright:

> [...] what is actually responsible, on this view, for the large degree of coherence and communicative success which our use of color vocabulary enjoys? Indeed, what is the justification for continuing to think of the use of such expressions as governed by rule? Knowledge of appropriate rules was supposed to constitute linguistic competence. But it cannot do so if competent usage essentially has a coherence which, in Dummett's view, the rules lack. Dummett's response needs supplementing with an explanation of our communicative success with such vocabulary in which the idea of knowledge of inconsistent rules has an ineliminable part to play. For either such knowledge is still to be a basic ingredient in competence or we should drop the idea (Wright 1987, 212).

Directly following this, Wright makes a related point:

> That brings us to [. . .] a decisive objection to Dummett's response. I do not see how we can rest content with the idea that certain implicitly known semantic rules are incoherent when nobody's reaction, on being presented with the purported demonstration of the inconsistency, i.e. the paradox – even if they can find no fault with it – is to lose confidence in the unique propriety of the response – e.g. "That's orange" – which the demonstration seems to confound (Wright 1987, 213).

I agree with Wright. We should try to account for the apparent communicative success of the use of vague vocabulary. Most of natural language lexical items are vague, and to dismiss the use as governed by inconsistent rules is bad theory. On the other hand, it seems to me absurd to try to rescue every single sequence of applications of vague predicates. Speakers do contradict themselves. It cannot be the goal of semantic theory to represent natural language as a foolproof means of making good sense (after all, people do paint themselves into corners and cut off the branches they are sitting on).

These considerations suggest that we try to find a semantic account of ordinary applications of vague predicates that does not reject all tolerance principles as unacceptable. For a forced-march sorites is not a sequence of ordinary applications. I shall therefore propose a combination of two strategies. For *ordinary* contexts of use, I propose that we opt for strategy a): reject the validity of the argument. For certain extreme contexts, on the other hand, I propose alternative d): the use of the vague predicate is incoherent.

In the next section, I shall sketch such an account, based on the observations of the first section. It is a contextual account, but unlike the mainstream of current contextual accounts, the most important feature will not be the *shift* between contexts, but an extra element in the context itself.

4 Central gap semantics

I think that the following correctly describes ordinary speaker psychology of ordinary applications of vague predicates: A predicate, like 'tall', is applied to a clear case of tallness or non-tallness, and although the speaker does not think that a sufficiently small difference, like 1 mm, can decide between being the one or the other, the difficult intermediate cases are simply "forgotten", or "ignored". It is enough that the case considered is clear, and that there are clear contrastive cases. There is no need to consider intermediate cases. We could say that intermediate cases are *dismissed*.

It is this dismissal of intermediate cases that can be modeled by means of interpretation that introduces significant central gaps in the domain of discourse in those cases where it is needed.

It is widely agreed that for interpreting (normal) utterances of sentences like

(14) Everyone went to bed at midnight

a contextually induced restriction on the quantifier domain is needed. The present proposal is that we extend the tool of quantifier domain restriction to give a *context semantics* for vague expressions that respects tolerance, bivalence as well as consistency for normal use. I shall propose that for a vague predicate F, like '. . . is tall', a speaker that accepts it as tolerant with respect to a particular dimension, uses it with a tacit assumption that *there is a restriction* on the domain of discourse, such that the domain in question has a significant central F-gap.

In this application, it is convenient to think of a central gap for a predicate F in a context c, relative to a dimension of variation, as a pair of real numbers (i, j) with respect to a measurement scale and a measure function $\mathcal{H}_F$. With an initially given domain D of individuals, the gap determines a proper subset of D, the set $F_c = \{a \in D : j \leq \mathcal{H}_F(a) \leq i\}$ of individuals in D whose F measures are in the gap. The gap then forces a *cut* in the domain, consisting in subtracting F_c from D. That is, the restricted domain is $D - F_c$.[8]

The type of semantics that I suggest is a context semantics.[9] It involves the assumption that in each context c, for each vague predicate F that is used in c, a tolerance level and a central gap is determined for F. It assumes the existence of a general *gap function* $\mathcal{G}$ that maps contexts of utterance on central gaps. Since the pair of number selected is arbitrary within limits, $\mathcal{G}$ must be a choice function with restrictions on the values it can give.

It must be required that in each context a full semantics is given 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 heads the list of restrictions, ordered 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.

[8]Some predicate are associated with more than one dimension of variation, e.g. predicates formed from simple predicate by means of connectives. In those cases we will need n-tuples of simple real number pairs, one for each dimension of variation.

[9]The account is more fully worked out in my *Vagueness and Domain Restriction*, to appear in a volume on vagueness and language use edited by Paul Egré and Nathan Klinedinst.

(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 of the truth of individual judgments. But this is a matter of further investigation.

That the gap must be at least equal to the tolerance level is necessary to preserve consistency in normal situations. If we make the further decision to let the size of the central gap for a predicate in a context be *equal* to the tolerance level for that predicate in the context, then the central gap can be regarded as determined by two well-known contextual standards: a standard of comparison and standard of precision.[10] We can simply identify the standard of precision with the tolerance level. If we further let the standard of comparison correspond to the center of the gap, then the gap is determined as a function of the standard of comparison and the standard of precision: where i is the value of the standard of comparison and k is the standard of precision, the gap is simply $(i + k/2, i - k/2)$.

There are then two basic ideas for the semantics: The first is that for each predicate F for which a gap is introduced, the extension of F consists of the individuals a in $D - F_c$ with measures above the gap ($\mathcal{H}_F(a) > i$) and the anti-extension of individuals b with measure below that gap ($\mathcal{H}_F(b) < j$). The second idea is that quantifiers are domain-restricted by means of the cut.

We can then verify that a tolerance principle, stated with binary quantifiers, such as

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

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's assume here that 'man' is non-vague. Then the antecedent of (15) 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 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 (15) is true.[11]

[10] Cf Lewis 1979, 244-46.

[11] The present account has some similarity to the account in Manor 2006, although the two were developed independently. Manor provides a non-standard semantics where the usual inductive premise fails if there is a suitable gap in the sequence of measures. According to Manor, the existence of a unique gap in the contextual domain provides a *natural demarcation* of the extension of a vague predicate in that context. To effect such a demarcation, the gap must be unique and well placed, and on Manor's account vague terms are used with the *presupposition* that there is such a gap. The present account differs from Manor's *inter alia* in that gaps are provided as part of interpretation, rather than declaring a speech act as failed when the presupposition isn't met.

5 Meaning and content

You might think that no adult human being, at least these days, who is 150 cm in height, could reasonably be counted as tall. That is, it would be part of the linguistic *meaning* of 'tall' that no interpretation function with an associated upper gap boundary below 150 cm is admissible.

A consequence of this view is that if there is no sharp boundary between admissible and non-admissible gaps, there will be no sharp boundary between admissible and non-admissible gap functions, and thus it would be unclear how a speaker could be interpreted. This is not, I think, a severe problem. We cannot anyway reasonably hope that our entire meta-language vocabulary is precise, and there is no good prior reason, from the present perspective, to believe that the domain of admissible gap functions, in some particular context of utterance, is sharply delimited. A leading idea of the present approach is that it is enough for adequate interpretation that there is a least one clearly admissible gap function (see below).

A more unwelcome result, however, is that two speakers A and B who disagree about how short an adult human can be and still count as tall, by such a view would be using 'tall' with different linguistic meanings. That is, they would *not* disagree substantially about the lower boundary of tallness, but would be speaking different languages, with phonetically and orthographically identical but semantically distinct predicates '. . . is tall'.[12] There would then be neither agreement nor disagreement between them on matters of tallness.

I find that implausible. Rather, with one proviso, if two speakers agree on the number uniformity principle (U) (stating so-called penumbral connections) of the predicate in question, then they share the concept. If they agree on the uniformity principle, then they agree on the dimension of comparison, and they agree on the direction relevant for the predicate (e.g. if taller, then more disposed to be counted as tall). The proviso is that for something to have a property as depending on its position in some order, then it must be possible that something has a lower (or, depending on the predicate, higher) position in that order. So, an object x is not tall unless it is possible that some other object y is shorter, and is not non-bald unless it is possible that some other object y has less hair. This means that for agreeing on the meaning of the predicate, two speakers must also agree on such absolute limits. Accordingly, I take it e.g. to be part of the meaning of 'bald' that a person with no hairs on his scalp is bald, but for no number greater than zero is it part of the meaning of 'bald' that a person with that number of hairs is bald.

Accordingly, I take two speakers who assign the same meaning in this respect to a tolerant predicate to agree on the general concept, such as the

[12]For an epistemicist, like Williamson (1994, 205-12), there is a sharp boundary of tallness determined for the language of the speech community, and from this perspective at least one of the two speakers would be mistaken about what the lowest admissible boundary is. Moreover, on this view, both speaker would be mistaken in thinking that there is a range of admissible alternatives. As many other non-epistemicists, I find such a determination of sharp boundaries from non-uniform use implausible.

concept of tallness, or baldness or heapness. However, such a concept is not individuated directly by application conditions. That is, the concept of being tall, for instance, is not individuated by a set of conditions C such that if an object x satisfies the conditions in C, x is tall and if it does not, it is not. The sorites paradox itself is a reason against this view, provided the premise of tolerance is accepted. For then, if there is such a concept of tallness, it is tolerant, and hence a tolerance principle is valid for that concept, which together with facts about the distribution of heights among adults of the world, leads to a sorites-type contradiction. We cannot avoid such a contradiction by changing the semantics of the concept, for concepts have their semantics built-in. Only for linguistic expressions, or other sign-like entities, can we devise alternative semantic theories. The conceptual semantic paradox can be avoided only if such a concept of tallness does not exist.

Rather, the concept of tallness, on the present view, is to be seen as a function from *standards of application* to extensions. And, on the present approach, admissible standards of application involve a contextual significant central gap. A central gap together with a measure function and the condition of belonging to the extension just in case one's measure is at least as great as the upper boundary, does fix the extension (and correspondingly for the lower boundary and the anti-extension).

For the issue of utterance and belief content, we must switch to an intensional framework. Within a possible-worlds framework, it is natural to regard the concept associated with a tolerant predicate as a function from possible worlds and standards of application to extensions. If we take the standard of application as the first argument, the value of the function is an ordinary intension, i.e. a function from possible worlds to extensions. For instance, an upper gap boundary for tallness of 180 cm determines an intension that for any possible world $\boldsymbol{w}$ fixes an extension consisting of the set of (adult male) humans in $\boldsymbol{w}$ that have a height of at least 180 cm.

It is not reasonable, however, to attribute such a precise belief content to the normal speaker. More plausibly, for each speaker and context there is a range of standards that the speaker is prepared to count as admissible. And plausibly, that range is not itself sharply delimited. That is, we have an unsharply bounded set of propositions that are admissible as intensions of the sentence in the context of utterance. This provides one dimension along which the expression 'believes that' is itself tolerant, and the present suggestion is, accordingly, that central gap domain restriction is to be applied to sentences containing it.

It is in line with this suggestion to take two speakers A and B to *agree* on a particular statement in a context c, not just conceptually but also doxastically, just in case there is a standard of application that is admissible to both A and B. This entails that for each tolerant predicate F involved in c, there is a pair (i, j) which is admissible for both $\mathcal{G}^A_F$ and for $\mathcal{G}^B_F$. As a consequence, doxastic agreement in a context is not transitive: there may a standard that is admissible to both A and B, and a standard that is admissible to both B and C but still no standard that is admissible to both A and C. But note that if two values are both admissible, a switch between them will not affect the truth value distribution over sentences. Hence, if A and B, and B and C, respectively, agree doxastically, they agree on truth values, and hence so do A and C, for the given linguistic fragment and the given domain. The non-agreement between A and C is conceptual, not a doxastic *dis*-agreement.

6 Problems

The basic strategy of the current proposal is to account for ordinary non-paradoxical use of vague expressions in one way, and allow inconsistency in the extreme cases where we have to do with sorites sequences. Accordingly, it is a necessary condition for the viability of the proposal that the account does give intuitively correct results for ordinary utterances, and in particular that ordinary utterances that do seem intuitively true or intuitively false, do not come out as incoherent on the proposed semantics.

It is, however, not obvious that this is the case. Potential sources of trouble are the combining of tolerant predicates, and the use of quantifiers. Here, there is not space for a thorough investigation of the matter. I can only discuss a few examples.

One thing that can happen is that objects are added consecutively to the domain in way that eventually eliminates the possibility of a central gap. If we say

(16) Julia is tall. So is Georgina, and so is Elsa, and so is Amanda ...

we may in the end populate the domain of quantification so that a sorites sequence results, given reasonable tolerance levels for ‘tall’. I don’t think this is a problem, however. Rather, this is one way in which the use of vague predicates can run into trouble.[13] This is perfectly in line with the present proposal.

The combination of tolerant predicates in a single context of utterance can give rise to the opposite effect. Since we need a gap and hence a domain cut for each vague predicate, the result might be that so much is cumulatively taken out of the domain of discourse that the topic is distorted. This consequence is avoided in case we are allowed context shifts that need not be conservative, in the sense that new cuts can be made and old cuts undone. This is an issue that requires further investigation.

But there is a problem even with a small cut. Suppose a basketball coach says

(17) Every player in my team is tall.

and it so happens that, given what counts as tall in the context, and given the contextual tolerance level, there is a sorites sequence in the team between the tall and the non-tall. Then some domain cut is needed, let’s say with the effect

[13] The idea that verb phrase ellipsis provides problems for other contextualist accounts of vagueness is due to Jason Stanley (2003).

that two players must go out of the contextually restricted domain. Have we thereby not *misrepresented* the content of the utterance of the coach? After all, the coach seems to be saying that *every player* of his team is tall, not that every member in a set consisting of all *except two* of the players in his team are tall.

The content ascribed to the utterance is not identical to the intuitive content, a content which would be incoherently ascribed. Rather, the content ascribed only *approximates* the intuitive content. The general idea of imposing contextual domain restrictions on utterance interpretation is that of getting the utterance content *right*. Isn't it used here with the opposite effect?

Let's note, first, that the intuitive truth value assignment is not affected. Intuitively, the utterance is false, since the team contains at least one non-tall player (or otherwise there would be no sorites sequence). But with a central gap, the cut in the domain is such as to leave anti-extension members in the domain. Hence, even with respect to the restricted domain the utterance comes out as false. Had all members been tall by the given standard, then no cut had been needed in the first place. So, either way, the intuitive truth value is preserved.[14]

Second: We have assumed that the utterance of the coach in itself makes sense. What doesn't make sense is to interpret it with the contextual tolerance level and the initially determined domain, i.e. the entire team. To bring out what is intuitively right, or in this case intuitively wrong, with the utterance, we need to deviate from the intuitive content ascription. This could be done in other ways, e.g. by assigning greater precision, i.e. a lower tolerance or even a zero tolerance level, going beyond any discrimination that the speaker himself would be prepared to make. This too would, on the current assumption, amount to distorting the intuitive content. Applying it across the board would have the effect of making the tolerance principle (15) come out false, despite being held true by the speaker, and thus force an error theory about the speaker, and about natural speakers in general to the extent that such tolerance principles are generally affirmed.

In order to assign content in a consistent way, some approximation has to be made. The type of approximation suggested here generally saves the intuitively assigned truth values and avoids the need of adopting an error theory about tolerance itself.

Department of Philosophy
Stockholm University

[14]We can indeed get more complicated cases with other quantifiers or determiners, like 'most', but there is not enough space here to discuss these cases. Some of them are more problematic, but then again problematic also on all standard accounts of vagueness.

References

Dummett, M., 1975, 'Wang's paradox', *Synthese* 30:301–24, . Reprinted in Dummett 1978. Page reference to the reprint.

Dummett, M., 1978, *Truth and Other Enigmas*, Harvard University Press, Cambridge, MA.

Fine, K., 1975, 'Vagueness, truth and logic', *Synthese* 30:265–300.

Graff, D., 2000, 'Shifting sands: an interest-relative theory of vagueness', *Philosophical Topics* 28:45–81.

Horgan, T., 1994, 'Robust vagueness and the forced-march sorites paradox', in J. E. Tombelin (ed.), *Logic and Language*, volume 8 of *Philosophical Perspectives*, Ridgeview, Atascadero, CA.

Keefe, R., 2000, *Theories of Vagueness*, Cambridge University Press, Cambridge.

Keefe, R. and Smith, P. (eds.), 1996a, *Vagueness: A Reader*, MIT Press, Cambridege, MA.

Keefe, R. and Smith, P., 1996b, 'Introduction: theories of vagueness', in R. Keefe and P. Smith (eds.), *Introduction: theories of vagueness*, 1–57, MIT Press.

Lewis, D., 1979, 'Scorekeeping in a language game', *Journal of Philosophical Logic* 8:339–59, . Reprinted in Lewis 1983. Page reference to the reprint.

Lewis, D., 1983, *Philosophical Papers. Volume I*, Oxford University Press, Oxford.

Machina, K. F., 1976, 'Truth, belief, and vagueness', *Journal of Philosophical Logic* 5:47–78.

Manor, R., 2006, 'Solving the heap', *Synthese* 153:171–86.

Raffman, D., 1994, 'Vagueness without paradox', *Philosophical Review* 103:41–74.

Raffman, D., 1996, 'Vagueness and context-relativity', *Philosophical Studies* 81:175–92.

Shapiro, S., 2003, 'Vagueness and conversation', in J. C. Beall (ed.), *Liars and Heaps*, 39–72, Oxford University Press, Oxford.

Shapiro, S., 2006, *Vagueness in Context*, Oxford University Press, Oxford.

Soames, S., 1999, *Understanding Truth*, Oxford University Press, Oxford.

Stanley, J., 2003, 'Context, interest relativity and the sorites', *Analysis* 63:269–80.

Williamson, T., 1994, *Vagueness*, Routledge, London.

Wright, C., 1976, 'Language-mastery and the sorites paradox', in G. Evans and J. McDowell (eds.), *Truth and Meaning*, 223–47, Clarendon Press, Oxford.

Wright, C., 1987, 'Further reflections on the Sorites paradox', *Philosophical Topics* 15:227–290. Reprinted in Keefe and Smith 1996a (with omission of section 5), 204–50. Page references to the reprint.