

Vagueness: why do we believe in tolerance?*

Paul Égré

Abstract

The tolerance principle, the idea that vague predicates are insensitive to sufficiently small changes, remains the main bone of contention between theories of vagueness. In this paper I examine three sources behind our ordinary belief in the tolerance principle, to establish whether any of them might give us a good reason to revise classical logic. First, I compare our understanding of tolerance for precise predicates and for vague predicates. While tolerance for precise predicates results from *approximation*, tolerance for vague predicates appears to originate from two more specific sources: *semantic indeterminacy* on the one hand, and *epistemic indiscriminability* on the other. Both give us good and coherent grounds to revise classical logic. Epistemic indiscriminability, it is argued, may be more fundamental than semantic indeterminacy to justify the intuition that vague predicates are tolerant.

1 Introduction

Logic thrives where paradox threatens. No wonder that vagueness and the sorites paradox have elicited so much work in the forty years that have passed since the birth of the JPL. With this anniversary, we may seize the opportunity to celebrate even more than that: a century of exciting logical and philosophical theorizing about vagueness. Let us look back to the beginnings: at the time Frege, Peirce and Russell invented modern logic, they intended it to serve a scientific ideal of precision and exactness. But they all warned us that two-valued logic might not be suitable for vague concepts as a result. In his attempt to found the mathematical principle of induction, Frege pointed out that the concept of a “heap of beans” is generally not hereditary in the sequence determined by the relation of differing by just one bean, “due to the indeterminacy of the concept of ‘heap’ ” (Frege, 1879). Peirce wrote to James that “it is a serious defect in existing logic that it takes no heed of the limit between two realms” (Peirce to James 1909, cited

(Revisions, July 2014). I would like to thank Frank Veltman and Jeff Horty for their invitation to contribute this paper. I am also indebted to Pablo Cobreros and Frank Veltman for comments, to Heather Burnett and to Laurence Goldstein and his students for discussions, and to audiences and colleagues at seminars held in Paris and NYU. Thanks to the European Research Council under the European Community’s Seventh Framework Program (FP7/2007-2013), to the project ‘Borderlineness and Tolerance’ (FFI2010-16984), Ministerio de Economía y Competitividad, Government of Spain, and to grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL.

in (Rescher, 1969)). And Russell went as far as to say that the law of excluded middle might be “not true when symbols are vague” (Russell, 1923).

A century later, a vast number of frameworks and approaches have been proposed to answer the challenge posed by vagueness and to clarify those intuitions. Among those theories, one of the persistent bones of contention, probably the main, concerns the treatment of the major premise of the sorites argument. In the classic Heap argument, this is the inductive premise that if n grains of wheat do not make a heap, then neither do $n + 1$, and in the case of the Bald Man, the premise that if n hairs on one’s head do not make one bald, neither do $n - 1$ hairs. In a more abstract form, the premise says that if a vague predicate P applies to an object x , and x and y differ very little in respects relevant to the application of the predicate P , then P also applies to y . Since (Wright, 1976), the premise is referred to as the *tolerance principle*. Wright famously called a predicate P “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”.

There could hardly be more disagreement on any philosophical subject than there currently is about the semantic status of the tolerance principle. On the epistemic view of vagueness ((Sorensen, 2001), (Williamson, 1994)) the principle is simply false. It is also false, though less straightforwardly, according to the supervaluationist definition of truth ((Fine, 1975)). In standard fuzzy logic, the principle is almost true, but it is not perfectly true (see (Williamson, 1994) for an exposition, and (Smith, 2008) for a comprehensive account). On some indeterminacy views of the nature of vagueness, the principle is neither true nor false (viz. (Schiffer, 2003), (Field, 2010)). On one family of contextualist approaches ((Kamp, 1981), (Raffman, 1994), (Fara, 2000), (Shapiro, 2006)), the principle is true as restricted to pairwise comparison, but not without this contextual restriction. On another family ((Gaifman, 2010), (Pagin, 2011), (Gómez-Torrente, 2011)), tolerance is a more fundamental desideratum, but the principle is only contingently sound, depending on whether some domain restrictions are in play.¹ Few theories treat (at least some version of) the principle of tolerance as valid unconditionally, although some do, including paraconsistent theories ((Hyde, 1997), (Ripley, 2013), (Weber, 2010)), and more recently, nontransitive logics ((Zardini, 2008), (van Rooij, 2012), (Cobreros et al., 2012b), (Cobreros et al., 2012a)). The reason so much disagreement persists among theories is not surprising: the tolerance principle is what induces paradox in a classical setting, which immediately casts doubt on its soundness. On the other hand, the notion of tolerance is often viewed as a defining feature of vague predicates as opposed to non-vague predicates (viz. (Dummett, 1975), (Wright, 1976), (Kamp, 1981), (Fara, 2000), (Greenough, 2003), (Eklund, 2005), (Gaifman, 2010), (Weber and Colyvan, 2011)).

Following a distinction made by (Fara, 2000), we should carefully distinguish three questions regarding the principle of tolerance. One is the psychological question: why are we so inclined to believe in the tolerance principle, even as we see that it leads to paradox?

¹When those restrictions would fail, some of those accounts accept the Dummettian view that vague predicates are simply inconsistent, see (Dummett, 1975) and (Pagin, 2011), (Gómez-Torrente, 2011).

Another is the epistemological question: are we founded to maintain our belief in the tolerance principle, despite the fact that it leads to paradox? And a third, the semantic question, is: granting the plausibility of the principle, should we seek to revise the laws of classical logic in order to validate it? These three questions are importantly different.² The main question this paper aims to clarify is the second question, the epistemological question. This issue, importantly, also touches on the psychology of vagueness, although I mean to interpret the question in a normative sense, such as: what are our best reasons to believe that the tolerance principle might be correct? Obviously, an answer to the epistemological question can be taken as a precondition to an answer to the semantic question, of whether tolerance may vindicate a revision of the laws of logic. If it turns out that none of our best reasons to believe in tolerance is solid enough, then a revision of logic is inappropriate.

The proposal I want to make in this paper is that we should distinguish at least three different sources behind our intuitions of tolerance. One source is not specific to vague predicates, it concerns the link between the notions of measurement, granularity and approximation, and it is relevant for the application of precise predicates. Basically, this notion of tolerance concerns the limited precision with which quantities can be measured. The idea is: so long as quantitative predicates are evaluated or reported with some approximation, there is some wiggle room in our use of them. In the case of vague predicates, by contrast, our intuitions of tolerance appear to originate from two further sources: one source is semantic indeterminacy, which supports the idea that vague predicates do not draw sharp boundaries but rather have an extended region of borderline cases. Another source is epistemic indiscriminability, which prevents us from drawing reliable distinctions between similar objects. Both those sources, I will argue, give us a good and congruent articulation of the notion of tolerance, in the form of gap principles. Such principles do actually give us a coherent motivation to revise logic, provided such principles are seen as expressing constraints on assertion, rather than as modalising the content of what is asserted.

2 Tolerance for precise predicates

Underlying the tolerance principle is the intuition that a slight difference along some dimension of comparison should not make a difference along a distinct, coarser-grained dimension. Folk psychology appears to be of two minds about this idea. One of the one hand, we find such universal sayings as: “that’s the straw that breaks the camel’s back” (in French: “c’est la goutte d’eau qui fait déborder le vase” – ‘it’s the water drop that makes the glass overflow’).³ On the other hand, we do have the intuition that, in many circumstances, a small difference will not make a difference. For instance, using an example briefly considered in (Hobbs, 2000), we typically believe that “adding a stamp to a letter does not change its weight enough for more postage to be required”.

²Fara’s articulation of the epistemological question differs somewhat from the present one, but this does not bear on the main distinction at stake here.

³See (Fine, 1975), (Gaifman, 2010), (Weber and Colyvan, 2011) on the consideration of such examples.

Part of the reason for this conflict of intuitions originates in measurement-theoretic considerations about the relation between the local or incremental level and the global or categorical level. The relation between the level of small increments and the level of categorical change is a relation between two scales with distinct granularities. For example, the non-tolerance intuition underlying the camel’s back proverb corresponds quite directly to the *Archimedean* property of the real numbers: take a fixed quantity A (like the maximum charge the camel’s back would theoretically support), and take an arbitrarily smaller but positive quantity ε (like the weight of some straw), there is an integer n such that $n\varepsilon$ will outweigh A .⁴

Our reasons to believe that adding a stamp to a letter will not affect postage are much less clear. One intuition people sometimes have about tolerance relates to the abstract property of *Density* shared by the rational numbers and the real numbers. The idea may be put as follows: consider a fixed weight A again, and a lighter weight A' such that the difference between them is a small quantity ε . For example, assume that A is the camel’s maximum charge, and A' the weight that you obtain by subtracting the weight ε of a given straw. Density says that you can find a lighter straw, of weight $\varepsilon/2$, such that the addition of this quantity will not outweigh A . We could keep adding straws to the camel’s back without ever breaking its back, provided the weights form a series of decreasing quantities converging to less than $(A - A')$.

Density, however, cannot be the right model for the letter and stamp. Obviously, a stamp is not an object whose weight can be arbitrarily decreased. Density, more generally, cannot be the right model for tolerance. Let x and ε range over weights with P the relevant predicate (“not breaking the camel’s back”). Density basically implies that:

$$(1) \quad \forall x \exists \varepsilon (Px \rightarrow P(x + \varepsilon))$$

In Wright’s definition of tolerance, quantifiers appear in the reverse order (“*there is* some positive degree of change... insufficient *ever...*”). That is, for P to be tolerant, it must be the case that:

$$(2) \quad \exists \varepsilon \forall x (Px \rightarrow P(x + \varepsilon)).$$

As (Pagin, 2011) writes, the density idea matches a weak and purely ordinal notion of tolerance. The notion intended by Wright’s definition is at least that of an interval scale, where increments are kept constant.

A different and more adequate account of the letter and stamp is in terms of the notions of approximation and granularity (see (Sauerland and Stateva, 2011)). In France in 2013, for example, postage varies as follows (for just the first two levels along the postal chart): up to 20g, postage for letters costs 63 cents. Between 20 and 50g, postage costs

⁴Regarding the French version of the proverb, we are certainly ready to concede that a given glass recipient has a determinate volume. The English version of the proverb is more interesting in a way, because there is arguably no absolutely determinate load that counts as the maximum load the camel can support. The fact remains, however, that whatever the finite range in which this maximum would fall, its upper bound can be outweighed by any small quantity.

1.05 euros. I have no idea of how much a standard stamp weighs, but we may assume that 100 stamps together do not weigh more than 10g, so that a stamp weighs at most 0.1 g. Consider the following sentence:

- (3) If a letter requires no more than 0.63 euros of postage, then sticking one stamp on it won't affect postage

Is this sentence true or false? Everyone would agree, I think, that the answer depends on the method and on the precision with which letters' weights are estimated. Currently, in some post offices at least, you find digital scales that display weights only up to 1g. What happens if you add a stamp to some quantity weighed on the scales? I did a couple of experiments. In the first, I dropped a single stamp on the empty tray and the number displayed remained "0g". In the second, I dropped a pack of envelopes, and the scales displayed "58g"; I then added a stamp on top, the display remained "58g". As limited as this sample might be, it is very tempting to induct that whatever quantity is initially displayed by the scales, the addition of a stamp will not modify what it displayed.

Supposing this were the case, here is how a purely "as if" explanation could go. Let us assume postal scales to be perfectly calibrated scales that estimate quantities with more precision than they display. They display the result only up to 1g, adopting the following rule: for every integer n , if a weight is estimated to be in the interval $]n - 0.5, n + 0.5]$, they round the result to n .⁵ Thus, when you first drop a stamp, we assume the scales detects a positive weight, but less than 0.5g, and rounds down the result to 0. Similarly, if you first drop a letter, and it finds 20.768, it rounds it down to 20. Now, assume that when you add a stamp of weight ε , the scales registers $20 + \varepsilon$, based on the previous number displayed (rather than estimated). If that were the case, it should round the result to 20 and display that number. In other words, for all letters whose rounded weight is less or equal to 20, if you consider the result of adding the rounded value of the stamp to the rounded value of the letter, that weight is still displayed as less or equal to 20.

Consider the predicate P "requiring no more than 0.63 euros of postage". Based on our assumptions, this predicate applies to all quantities whose actual weight is in the real interval $[0, 20.5]$. For this predicate, and for any quantity ε strictly less than 0.5, we get the following generalization, where $\lfloor x \rfloor$ is the function that, to each real number x associate the value n whenever x is in the interval $]n - 0.5, n + 0.5]$:

$$(4) \quad \forall x \in \mathbb{R} (P[\lfloor x \rfloor] \rightarrow P(\lfloor x \rfloor + \lfloor \varepsilon \rfloor))$$

This generalization, however, does not predict what will happen if you first take a letter of weight 20.45g, weigh it, then remove it, reset the scales to 0, add a stamp to the letter it, and weigh it again. If ε exceeds 0.05, then $\lfloor x + \varepsilon \rfloor$ will reach 21. So, obviously, we do not have the generalization that:

⁵Standard rules differ, in that they round to the even round value, see (Kirkup and Frenkel, 2006). The specificity of the rounding rule chosen does not matter however in what follows, we may even consider nonsymmetric approximation intervals, and adapt our analysis accordingly.

$$(5) \quad \forall x \in \mathbb{R}(P[x] \rightarrow P[x + \varepsilon])$$

As we can see, (4) and (5) offer two distinct ways of making sense of (3). According to the first, the intuition of tolerance is valid, but not according to the second. The first, arguably, is implausible with respect to the way actual scales work (though not necessarily implausible as one way we might fancy approximation), for it assumes that in the sequential addition of weights, the weight displayed is the sum of individually rounded weights, rather than the rounded sum of the weights.

A related but different intuition one may have about the validity of (3) is not by reference to a particular weighing algorithm, but more directly, in terms of contextual domain restriction. The idea is the following: when considering the antecedent of (3), the choice of a particular scale of measurement may simply tell us to disregard letters whose weight is not a multiple of the unit interval. Arguably, the consideration of non-unit values becomes relevant only when reaching the consequent of the conditional, namely with mention of the stamp, whose weight can only be appreciated using a more fine-grained scales. Under those assumptions, the interpretation of (3) would be as follows:

$$(6) \quad \forall x \in \mathbb{N}(Px \rightarrow P(x + \varepsilon))$$

Like (4), this sentence is true under our assumptions, but its truth essentially depends on the granularity chosen. Let us write $\mathbb{N} \cdot 10^{-1}$ for the set of all multiples of 10^{-1} . Assume now that the scales were to display weights with a precision of 0.1g, using the same approximation rule. For example, any weight in the interval $]19.95, 20.05]$ is rounded to 20. At this more refined level, we can no longer be sure that adding a stamp will have no effect. If the precise predicate P (“requiring 0.63 euros of postage”) is now assumed to apply to all weights in the real interval $[0, 20.05]$, the following sentence is no longer true when ε can be 0.1g, since x might be equal to 20.0g:

$$(7) \quad \forall x \in \mathbb{N} \cdot 10^{-1}(Px \rightarrow P(x + \varepsilon))$$

Obviously, the finer the granularity of the scale, the more likely a fixed small weight ε is to make a difference. In particular, if we allow quantification over arbitrary real numbers, the finest granularity of all, then we are back to the Archimedean intuition we started with: any ε will make a difference. However, given an interval scale built on intervals of unit a , with the approximation rule that values in the interval $]n \cdot a - \frac{a}{2}, n \cdot a + \frac{a}{2}]$ are rounded to a , any quantity ε strictly less than $a/2$ is such as to make no difference. That is, for every numerical predicate P true of all and only numbers strictly less than some multiple of the unit interval a on the scale, we have:

$$(8) \quad \forall \varepsilon < \frac{a}{2} \forall x \in \mathbb{N} \cdot a (Px \rightarrow P(x + \varepsilon)).$$

The analysis in (8) applies quite well to the kind of examples discussed by Sauerland and Stateva concerning the relation between tolerance and approximation: their example is that if you date a particular dinosaur bone to be “1 million years old”, then three days later, although it is strictly speaking older, it would be inappropriate to date it to be “1 million years and three days”. Three days later, it is still appropriate to call it

“1 million years old”, because a difference of three days falls short of reaching the next relevant graduation on the appropriate scale. More generally, the pattern in (8) can be seen to correspond to what (Williamson, 1992) calls a *margin for error principle*. When we say that the bone is 1 million years, or consider letters weighing no more than 20g, we do it within a margin of error set by the approximation interval around each of those values. What (8) says is basically that if a precise predicate P applies to some quantity x within a given margin of error, P is still applicable to any quantity $x + \varepsilon$ provided ε does not exceed the margin.⁶

However, (8) is not the only conceivable way of understanding the letter example we started with, and it is not obvious that it can be generalized to all similar cases. To see this, let us consider a different example:

(9) If your bare foot fits in a shoe of size 8, then it will still fit if you put a sock on.

There is I believe a very natural way of understanding a sentence like this one (and similarly for (3)), namely that “*most of the time*, or *typically*, or *generally*, if your bare foot fits in a shoe of size 8, then it will still fit if you put a sock on”. Even if we know standard sizes for shoes to go by inches and half inches, I think we do not understand a sentence like (9) by restricting attention only to bare feet that are integral multiples of half inches. What we appear to mean, rather, is that it is unlikely that one’s foot could fit in a shoe of size 8, and not fit with a sock on. Similarly for (3), although we cannot rule out the possibility that, occasionally, the addition of a stamp will make the scales tip from 20 to 21, we understand that it is so unlikely as to be a negligible possibility. Note that if this is how we understand sentences like (3) or (9), this is sufficient to conclude that we do not ever perfectly believe in the tolerance principle for such precise numerical predicates. Rather, we believe in such conditional sentences with a high degree of probability.

How this probability should be determined is not completely straightforward. One robust intuition, however, is that the smaller the size of the increment, the higher the probability will be. (Edgington, 1997) proposes to view this probability as 1 minus the ratio of the measure of the increment ε to the measure of the predicate (that is, 1 minus the relative size of the interval where an increment would make a difference). A similar idea, interestingly, was put forward by (Borel, 1907) in what is probably the first extended treatment of the sorites paradox published in the 20th century (see (Egré and Barberousse, 2014) for a translation and detailed presentation). Borel considered the following problem: consider two monetary scales with distinct discrete granularities, a scale for wholesale price, where coinage can vary by increments of 1/2 cent, and a scale for retail price where coinage can vary by increments of 5 cents. What is the effect of one increment of 1/2 cent of wholesale price on retail price? Borel described as a fallacy the idea that *exactly* 10 diminutions of 1/2 cent in wholesale price are needed

⁶The analysis may also be related to the central gap account of tolerance defended by (Pagin, 2011), where the tolerance principle is pragmatically restricted by the assumption that individuals in the domain do not fall within a central gap whose size is at least that of a tolerance level fixed by the context for each vague predicate.

to reach a diminution of 5 cent in retail price. His idea, rather, was that the effect of an increment or decrement in wholesale price is better described by a random variable having a probability 1/10 of occasioning an increment or decrement of 5 cents in retail price. In his analysis too, 1/10 corresponds to the ratio of wholesale price unit to retail price unit.

Summarizing what we said so far, we see that, as applied to precise numerical predicates, our intuitions about the validity of the tolerance principle oscillate. On the one hand, we can have full credence in at least one interpretation of (3), involving multiple domain restrictions in order to ensure its validity. On the other hand, when such restrictions are no longer in play, we appear to only have a high credence in such tolerance conditionals, dependent on the relative measure of the increment to the measure of the predicate. Either way, we seem to have reached a dilemma: either we are in a position to fully believe in tolerance, because the principle comes out as weaker than it appears to be; or we place no restriction on the tolerance principle, and we can only believe it to a high but imperfect degree.

3 Tolerance for vague predicates

In the previous section we have focused on what the notion of tolerance might mean for precise predicates, namely predicates with a determinate and sharp boundary along some predefined scale of measurement. What we have found, basically, is that for such predicates, the idea of tolerance is tied to the notion of approximation, and that it concerns the limited precision with which we can measure or report theoretically precise quantities. The question that now awaits us is: what about the idea of tolerance for vague predicates? How should the idea be articulated, and does it have the same foundation, or a distinct foundation?

The answer depends on whether one thinks vague predicates and precise predicates are essentially distinct. On the epistemic view, for example, vague predicates are not fundamentally distinct from precise predicates, they come with a sharp but unknowable boundary, and they are in fact subject to the same constraints. For Williamson, we are simply confused about the tolerance principle: when we believe a predicate like “thin” to be tolerant, what we rationally intend is actually the margin for error principle (if x is known to be thin within a margin of α , and x and y do not differ by more than α , then y is thin), but not the tolerance principle proper (if x is thin, and x and y differ by no more than α , then y is thin).⁷

Appealing though this answer might be, there is a widely shared intuition that there may be more to the idea of tolerance for vague predicates than expressed by the margin for error principle. On the opposite view, tolerance is part and parcel of the meaning

⁷See (Williamson, 1992): “Fortunately, ‘thin’ is not governed by the tolerance principle; it is governed by the margin for error principle (!), which generates no sorites paradox”. Williamson’s margin of error principle implies that we cannot know where the cutoff lies. According to Williamson, it is a verificationist bias that urges us, from the impossibility to know Px and $\neg Py$ of any x and y such that $x \sim_P y$, to believe that there are no such x and y .

of vague predicates (see (Dummett, 1975), (Eklund, 2005), (Gaifman, 2010) for particularly clear statements of such a view). Importantly, and even to those who defend the view, this does not imply that a vague predicate should necessarily satisfy the tolerance principle in *all* contexts of use. Depending on the context, granularity considerations do interfere with the sorites-susceptibility of vague predicates. (Gaifman, 2010) points out that the vague predicate “large number of fingers”, as applied to someone’s hands, may be such that $n + 1$ is a large number of fingers, but n is not. An even clearer example is due to (Burnett, 2012) who points out that so-called penny candies, which cost a penny, are definitely not expensive, but would be judged expensive if they were sold 2 cents or more. But this does mean that *expensive* should be considered a nonvague adjective generally speaking. As argued by Burnett in her work on relative adjectives, what appears to matter for a predicate being vague is the *existence* of a context in which the granularity is such as to make the predicate sorites-susceptible. Burnett calls this feature *potential vagueness* and rejects the idea, quite rightly in my opinion, that for a predicate to be vague, it should be universally tolerant, as opposed to merely existentially tolerant. In the case of Gaifman’s example, we therefore ought to distinguish the potential vagueness of “large” from the contribution of “number of fingers”.

The question then is what could justify the idea that tolerance, given an appropriate context and granularity, is constitutive of the meaning of vague predicates. To gain some insight into the problem, I propose to consider (Borel, 1907)’s pioneering account as one illustration of that view. Borel was clearly of the opinion that the principle of tolerance cannot be fully endorsed in general, not only on pain of contradiction, but also, as explained above, because he saw the effect of slight modifications of the respects relevant for the application of a predicate as necessarily affecting the probability of subsuming the object under the predicate. Nevertheless, Borel opened his reflections on the heap paradox with what may be viewed, by way of anticipation, as a frontal anti-epistemicist claim regarding vague predicates:

“One grain of wheat does not make a heap; neither do two grains; nor do three grains; ... on the other hand, if a million grains are gathered, we can agree that they constitute a heap. What is the exact limit? Should we say that 2342 grains, for instance, do not make a heap, whereas 2343 grains do? This is obviously ridiculous.”

Borel’s target in this example is stronger than the existential denial of tolerance: it amounts to stating a particular instance of it. Because of that, though Williamson endorses the existential sentence resulting from the negation of tolerance, he might agree with Borel that we should not say of any particular value n that n grains make a heap and $n + 1$ do not, in so far as: i) it follows from the margin of error principle that it is impossible for us to know that n grains do not and $n + 1$ do make a heap, and ii) we should not assert what we do not know. However, we may take Borel’s question to challenge even the view on which, supposing we knew all the numerical facts about collections of grains, and all actual and counterfactual dispositions of speakers to assert

heapdom of such collections, there would be a determinate value n of which we should assert $P(n + 1)$ and Pn .

How could Borel be so confident in the idea that negating the major premise of the sorites is absurd? One answer we find sketched in Borel’s paper concerns the link between vagueness and the notion of *semantic indeterminacy*. For Borel, there is a “necessary indeterminacy of verbal definitions”, whose main effect is that competent speakers can issue distinct yes/no verdicts regarding the application of a predicate to an object on distinct occasions of use, depending in part on their backgrounds and interests.⁸ This multiplicity of verdicts, which is the main basis of supervaluationism about vagueness, can indeed be used to articulate one notion of tolerance specific to vague predicates (see below).

But semantic indeterminacy is not the only alleged source for the tolerance principle in the case of vague predicates. Another widely discussed basis for the tolerance principle in the case of vague predicates is the notion of *epistemic indiscriminability*, namely our inability to reliably keep track of small differences.⁹ The idea can be put negatively, indiscriminability standing as a limit on our capacity to detect the limits of precise categories (Williamson, 1994), but often also positively, as a constitutive aspect of the process of category formation ((Raffman, 1994), (Fara, 2000), (Kennedy, 2011), where the positive articulation corresponds to the notion of salient similarity; see below).

Of those two phenomena, semantic indeterminacy, and epistemic indiscriminability, is one phenomenon derivable from the other, or are they independent? This question, I take it, remains one of the most important open questions regarding the nature of vagueness. I will not try to answer it directly here. What I would like to show first is that irrespective of whether semantic indeterminacy or epistemic indiscriminability is taken as more fundamental, both ideas converge on a common way of articulating the notion of tolerance at the logical level.

Let me consider the link between tolerance and semantic indeterminacy, first by returning to Borel’s claim. Although Borel does not make his point explicit, a natural

⁸See particularly (Wright, 1994) and (Raffman, 1994) on permissible disagreement in borderline cases. The resulting notion of indeterminacy fits what Smith calls *plurivaluationism*, what (Eklund, 2005) calls *second level indeterminacy*, or what (Shapiro, 2006) calls *open-texture* (in a sense compatible with but weaker than Waismann’s original sense), to characterize the compatibility of the meaning of an expression with a multiplicity of different verdicts. This notion of semantic indeterminacy is compatible with stronger notions, such as what Eklund calls first-order indeterminacy, for when a predicate is semantically gappy, but it is not mandated by it.

⁹The notion of indifference is sometimes judged even more fundamental than that of indiscriminability, for cases in which we are in a position to make a distinction, but in which the difference is judged practically irrelevant. Arguably, however, the notion of indiscriminability is more explanatory of the phenomenon of tolerance. (Gaifman, 2010) writes that a difference between 41 and 40 chimpanzees is as discriminable in principle as a difference between 5 and 6 chimpanzees, but practically less relevant when ascribing “large” to a community of chimpanzees. This is correct, but it would be less easily discriminable if presented perceptually. He also considers that “large” is not a perceptual predicate in that example, but this is controversial. Our mental representation of numerosities, relative to some given unit, and even for more clearly nonphenomenal predicates such as “rich” or “expensive”, may remain highly dependent on our perceptual ability to discriminate between quantities. See (Fulst, 2011) and (Feigenson et al., 2004) for arguments in that direction.

way of interpreting his “no sharp boundary claim” is as follows:

- (10) there is no number n of grains such that we should say that n makes a heap, and we should say that $n + 1$ does not make a heap

This normative generalization itself can be grounded in the following fact, assuming that “we should say that A ” is true if and only if A is determinately the case, where “ A is determinately the case” can be taken to mean that all competent speakers would agree that A :

- (11) There is no number n of grains such that it is determinately the case that n makes a heap and it is determinately the case that $n + 1$ does not make a heap.

Since (Wright, 1992), principles like (11) have been called *gap principles* (see (Fara, 2003), (Cobrerros, 2011)), since basically, what (11) requires is a gap between cases that are determinately P and cases that are determinately not P . Such principles have been criticized as conducive of paradox in relation to the phenomenon of higher-order vagueness (Fara, 2003), and looked at with suspicion for that matter, but mostly based on a rule specific to supervaluationism (the rule of D-introduction, according to which “Determinately A ” should follow from A). This inference is independently objectionable, however, and as I see it, it would be hasty to throw out gap principles in an account of first-order vagueness simply because they interact in bad ways with a rule that is dubious.¹⁰

When we look at the recent literature on vagueness, it is indeed a striking fact that we find so many congruent articulations of the notion of tolerance in terms of principles of this form, even outside of the supervaluationist tradition.¹¹ (Greenough, 2003)’s minimal theory of vagueness, for example, proposes the following notion of epistemic tolerance as generic for vagueness:

“there are no close cases in which it is known that a sentence takes a certain truth-state in one case and known that this sentence takes the complementary truth-state in the other close case”.

¹⁰See in particular (Cobrerros, 2011) for a defense of gap principles against D-introduction, and work in progress by P. Pagin opposing the same rule, based on the assumption that “determinately” operators behave as intensifiers. For a recent assessment of higher-order gap principles, see also (Zardini, 2013).

¹¹See among others (Wright, 1995), (McGee and McLaughlin, 1995), (Edgington, 1997), (Greenough, 2003), (Egré, 2009), (Lassiter, 2011) or (van Rooij, 2012). For example, (McGee and McLaughlin, 1995) write: “One existential claim we surely do *not* want to make is this: there is an n such that the n th tile definitely looks red and the $(n + 1)$ st tile definitely does not look red. Such a claim would allege a justifiable distinction where there is no significant difference”. Likewise, (Edgington, 1997) writes: “no single grain makes a difference between a heap and a non-heap’. If this means: no single grains makes a decisive difference – takes you from a clear heap to a clear non-heap – then it is true. If it means: no single grain at all makes any difference to heapdom, then it is false”. Even (Smith, 2008)’s defense of closeness, the idea that close cases should receive close semantic values, may be reanalyzed in terms of gaps, as the idea that the assignment of opposite semantic values to Pa and Pb implies that a and b should be sufficiently far. See (Egré, 2011) and (Cobrerros et al., 201x) for further discussion.

Significantly for our purpose, Greenough’s approach considers epistemic indiscriminability to be more fundamental to vagueness than semantic indeterminacy, but we see that his notion of tolerance is again a gap principle, only using knowledge instead of definiteness. Abstracting away from the difference between knowledge and determinacy, the logical form behind all gap principles is the following, where $x \sim_P y$ means that x and y are close cases in P -relevant respects (see (van Rooij, 2012), (Cobreros et al., 2012b)):

$$(12) \quad \forall x \forall y (x \sim_P y \rightarrow (\Box Px \rightarrow \neg \Box \neg Py))$$

Assuming a principle like (12) gives us the right description of the way we understand the idea of tolerance for vague predicates, two questions remain to be answered. The first is: does the validation of this principle motivate in any way a revisionary stance toward classical logic? The second is: should we refer gap principles primarily to the notion of semantic indeterminacy, or to the notion of epistemic indiscriminability?

Let me address the logical question first. *Prima facie*, the gap principle stated in (12) appears as a simple weakening of the original tolerance principle, since it basically involves a modalization in both the antecedent (strengthened) and consequent (weakened). If so, it may seem that the right logic for vagueness is just classical logic augmented with modal operators. Arguably, however, the modal operators present in (12) are not operators we necessarily want to refer to the content of sentences with vague predicates, but they are better seen as assertion modalities (remember, again, Borel’s quote: “should we *say* that...”). If we take seriously the idea that logic in general is a way of coding relations between our assertive commitments,¹² then the tolerance intuition expressed in the gap principle (12) arguably gives us a reason to modify the logic.

An illustration of this perspective on vagueness and logic can be found in (Cobreros et al., 2012b) and in (Cobreros et al., 2012a), where instead of viewing \Box as an object-language operator on a par with other ordinary predicates of the language, \Box is treated as an operator of the metalanguage, expressing a notion of strict or strong assertion (\models^s), with its dual, $\neg \Box \neg$, expressing a notion of tolerant or weak assertion (\models^t for $\not\models^s \neg$). The basic articulation we give of (12) is, given a suitable model M for vague predicates:

$$(13) \quad \text{for all } x \text{ and } y \text{ such that } x \sim_P y, \text{ it is not the case that: } M \models^s Px \text{ and } M \models^s \neg Px$$

That is, no two close cases in P -relevant respects are such that one can (strictly) assert predicate P of the one, and (strictly) deny P of the other. Equivalently, in a form immediately closer to the tolerance principle in conditional form, if two cases are close, then if P is assertible strictly of the first, it is assertible tolerantly of the second:

$$(14) \quad \text{for all } x \text{ and } y \text{ such that } x \sim_P y, \text{ if } M \models^s Px \text{ then } M \models^t Px$$

With the admission of such metalinguistic operators, the next question is: which semantics is adequate to make (14) a valid principle? There is, in fact, no unique answer

¹²See particularly (Zardini, 2008), (Ripley, 201x) and (Cobreros et al., 201x) for a defense of this inferentialist conception.

to this question (to put it otherwise, no unique solution to the equation or implicit definition laid out in (14) for s and t). (Zardini, 2008), the first to advocate a treatment of the tolerance principle in terms of a distinction between levels of goodness for assertibility, originally presented distinct logics compatible with the desideratum in (14). In (Cobreros et al., 2012b) and in (Cobreros et al., 2012a), similarly, we consider two further and mutually distinct answers to this question. Both those answers have in common the fact that in each of (Cobreros et al., 2012b) and (Cobreros et al., 2012a), \models^s and \models^t are treated as dual notions of truth for sentences, matched by dual consequence relations, and each giving rise to a distinguished nontransitive notion of mixed consequence (called st for “strict-tolerant” in the first framework, and pb in the second, for “super-sub”). In (Cobreros et al., 2012b), in particular, strict truth is equivalent to strong Kleene truth, and tolerant satisfaction with its dual, LP-truth, while in (Cobreros et al., 2012a), strict truth corresponds to supervaluationist truth (super-truth), and tolerant truth to subvaluationist truth (sub-truth). The choice between those two frameworks (or indeed other frameworks that would fit the tolerance pattern in (14)) is not dictated purely by (14), but it depends on further desiderata that one thinks a consequence relation should satisfy, and that one thinks a theory of vagueness should satisfy. For example, the s’valuationist framework of (Cobreros et al., 2012a) does not allow us to validate the tolerance principle in axiom form, although it allows us to validate each instance in rule form, whereas the TCS framework of (Cobreros et al., 2012b) allows us to validate tolerance both as a rule and as an axiom. Also, if we set aside logical consequence to consider what we is assertible of a borderline case, even clearer differences emerge: in st , a borderline case can be tolerantly asserted to be “ P and not P ”, but this is not so in the pb framework, or indeed in Zardini’s framework.

More to the point of our epistemological inquiry, in the TCS framework of (Cobreros et al., 2012b), the relation \sim_P of closeness between elements of the domain is, for vague predicates, a reflexive, symmetric and nontransitive relation taken to represent the notion of epistemic indiscriminability (similarity, or indifference) between two elements. In the s’valuationist framework of (Cobreros et al., 2012a), two individuals a and b are in the closeness relation \sim_P provided two (possibly identical) bivalent models (precisifications) give identical truth-values to Pa and Pb . The latter notion of closeness is, arguably, much weaker than the former, since it does not account for why and when verdicts would be identical for a given speaker. A good example of the difference between the two frameworks concerns, for example, the constraint on categorization discussed by (Raffman, 1994) and (Fara, 2000), called the constraint of *salient similarity*, which says that if you consider two individuals a and b , and they differ very little, then one should either accept Pa and Pb for both, or accept $\neg Pa$ and $\neg Pb$ for both. If this notion of acceptance is defined in terms of strict assertion again, then one way of representing it is directly in terms of model restriction (see (Egré, 2011)): if you restrict a model M to just the pair $\{a, b\}$ with $a \sim_P b$, then $M, \{a, b\} \models^s Pa$ iff $M, \{a, b\} \models^s Pb$, and likewise $M, \{a, b\} \models^s \neg Pa$ iff $M, \{a, b\} \models^s \neg Pb$. Both predictions are borne out in the indiscriminability based framework of (Cobreros et al., 2012b), but not in the supervaluationist framework of (Cobreros et al., 2012a) for the corresponding notion of

supertruth.¹³

Judging from that perspective, this difference goes in favor of a definition of closeness in terms of epistemic indiscriminability, rather than in terms of extrinsic identity of judgments. That, of course, is not enough to conclude that an understanding of tolerance based on epistemic indiscriminability is necessarily prior to an understanding of tolerance based on the notion of semantic indeterminacy (in the plurivaluationist sense of verdict multiplicity explained). Even in a supervalueationist framework, closeness need not be defined (and is not typically defined) in terms of within- or between-subject identity of verdicts for distinct objects. Consider again the classical heap paradox. In the standard supervalueationist treatment, the assumption is that two collections of grains are close if their number of grains differs very little (typically, by just one grain). However, even under that understanding of closeness, we need more than semantic indeterminacy (qua plurality of verdicts) to support the tolerance intuition expressed in (12). The reason why there is no number n such that it is super-true that $n + 1$ grains make a heap and supertrue that n grains do not make a heap is not just indeterminacy in the sense of semantic variation between speakers. What is also needed is: i) the *gap* assumption that Pa and $\neg Pb$ are determinate for some distant enough a and b , and ii) the assumption that all speakers abide by the *monotonicity* principle that if i grains make a heap, then so do j grains for all $j > i$.¹⁴ Without the latter assumption, and even assuming the gap assumption, it might happen that speakers who agree that 2000 grains don't make a heap but 2100 do, contingently agree that 2042 grains don't make a heap but that that 2043 do, and again diverge on 2044 and at other positions between 2000 and 2100. This means that in order to get the notion of tolerance expressed in (11) out of the idea of semantic indeterminacy, we need more than just the assumption of a plurality of verdicts in some region and the ordinary notion of closeness between cases, we crucially need the monotonicity assumption. By contrast, the requisite of epistemic tolerance expressed by Greenough follows directly, and without additional assumptions, from the relation of closeness assumed in TCS.

Can we conclude that epistemic indiscriminability between close cases is more fundamental than semantic indeterminacy in justifying our intuition that vague predicates are tolerant? This has some appearance of plausibility, but this should not be taken to imply that epistemic indiscriminability *on its own* can ground our intuitions of tolerance. Consider again the strict-tolerant account of vague predicates. What the account says is that inasmuch as one can assert strictly of an object that it is tall, then, given an individual that is indiscriminable in height, one cannot strictly deny of that object that it is tall. But of course, the account does not say that inasmuch as one can assert *tolerantly* of an object that it is tall, given an individual indiscriminable in height from it, one

¹³Incidentally, it would also fail in the trivalent reinterpretation of TCS presented for instance in (Cobreros et al., 201x), where closeness is defined directly in terms of distance between semantic values, rather than as a relation over objects proper. This feature is not essential to the trivalent account, however.

¹⁴Monotonicity is what (Fine, 1975) calls an internal penumbral connection. I take the term “monotonicity” from (Nouwen, 2011). It is used by (Gaifman, 2010) in the same sense. (Burnett, 2012) uses the term “scalarity”.

cannot *tolerantly* deny of the latter that it is tall (otherwise, there would be no way of explaining how we can switch category in borderline cases). Indiscriminability supports tolerance only in relation to strict assertion. In order to explain intuitions of tolerance, indiscriminability is therefore not sufficient: the notion of strict assertibility too must be explained. What does it mean, then, to say that a vague predicate is strictly assertible? We may say that a vague predicate is strictly assertible of an object whenever one can “anchor” one’s representation of the predicate to that object. Call such an object a *good case* of application of the predicate.¹⁵ A good case, in other words, is one for which one can issue a reliable judgment that the object exemplifies the predicate. A good case must therefore be a case for which indiscriminable changes in the relevant respects do not produce conflicting verdicts in the same subject (compatibly with good cases differing across subjects and contexts of use). Our belief in tolerance is fundamentally tied to indiscriminability relative to good cases; picture a borderline case of a tall man, we no longer believe that any man whose height differs ever so slightly will count as tall. This consideration comes close to Williamson’s claim that the true principle behind the tolerance principle is a margin of error principle. But the idea is not quite the same: on the present view, the principle of tolerance is not fundamentally tied to avoiding error around a predefined boundary. The difference may be put as follows: the view is not that in order to *know* an object of a given height to be tall, any height indiscriminable from it must also be tall. Rather, it is that to the extent that I can *coherently and justifiably* pick an object of a given height to anchor and deploy my representation of what counts as “tall”, then I can’t be charged of incoherence in predicating “tall” of an object whose size differs indiscriminably.¹⁶

4 Conclusion

We have identified three main sources behind our belief in the tolerance principle: approximation, semantic indeterminacy, and epistemic indiscriminability. Approximation accounts for one weak form of tolerance, which concerns the limited precision in our use of precise predicates. Semantic indeterminacy on the other hand gives us access to a distinct understanding of tolerance, this time pertaining to the lack of a sharp boundary between the determinate extension of a predicate and the determinate extension of its negation. Prima facie, the difference between vague and non-vague predicates may appear to lie in just that feature: even as we have fixed a comparison class, a predicate like “heavy”, unlike “weighs more than 20g”, does not specify an explicit and determinate threshold for objects to count as heavy. The existence of an area of competing verdicts may thus be viewed as the main source for our reluctance to assign a sharp cutoff to vague predicates. What we have argued, however, is that semantic indeterminacy may not be the ultimate source behind our intuition of tolerance in the case of vague predicates. The notion of epistemic indiscriminability, namely the impossibility

¹⁵Compare with (Zardini, 2008)’s talk of the “good” semantic values a sentence can have.

¹⁶See particularly (Raffman, 1994) and (Shapiro, 2006) for related views about tolerance and semantic competence.

to draw reliable boundaries between highly similar cases, is arguably more explanatory at the individual level. Importantly, however, the notion of epistemic indiscriminability is compatible with the fact that we come to draw lines for vague predicates on particular occasions of use. Because of that, a logic for tolerance is not a logic in which lines are prohibited, it is a logic in which no sharp line can separate cases which we can strictly assert to be P from cases which can strictly assert to be not P .

References

- Borel, E. (1907). Un paradoxe économique: le sophisme du tas de blé et les vérités statistiques. *La Revue du Mois*, 4:688–699. English translation by P. Égré and E. Gray, forthcoming in *Erkenntnis*.
- Burnett, H. (2012). *The Grammar of Tolerance: On Vagueness, Context-Sensitivity, and the Origin of Scale Structure*. University of California, Los Angeles. PhD thesis.
- Cobrerros, P. (2011). Supervaluationism and Fara’s argument concerning higher-order vagueness. In Egré, P. and Klinedinst, N., editors, *Vagueness and Language Use*. Palgrave Macmillan.
- Cobrerros, P., Egré, P., Ripley, D., and van Rooij, D. (201x). Vagueness, truth and permissive consequence. In D. Achourioti, H. Galinon, J. M., editor, *Truth*. Springer. Forthcoming.
- Cobrerros, P., Egré, P., Ripley, D., and van Rooij, R. (2012a). Tolerance and mixed consequence in the s’valuationist setting. *Studia Logica*, pages 1–23.
- Cobrerros, P., Egré, P., Ripley, D., and van Rooij, R. (2012b). Tolerant, classical, strict. *The Journal of Philosophical Logic*, pages 1–39.
- Dummett, M. (1975). Wang’s paradox. *Synthese*, 30(3-4):301–324.
- Edgington, D. (1997). Vagueness by degrees. In Keefe, R. and Smith, P., editors, *Vagueness: a Reader*, pages 294–316. MIT Press.
- Egré, P. (2009). Soritical series and Fisher series. In Hieke, A. and Leitgeb, H., editors, *Reduction: Between the Mind and the Brain*, pages 91–115. Ontos Verlag.
- Egré, P. (2011). Perceptual ambiguity and the sorites. In Nouwen, R., van Rooij, R., Sauerland, U., and Schmitz, H., editors, *Vagueness in Communication*, pages 64–90. Springer.
- Egré, P. and Barberousse, A. (2014). Borel on the heap. *Erkenntnis*. Forthcoming.
- Eklund, M. (2005). What vagueness consists in. *Philosophical Studies*, 125(1):27–60.
- Fara, D. (2000). Shifting sands: an interest-relative theory of vagueness. *Philosophical Topics*, 28(1):45–81. Originally published under the name “Delia Graff”.

- Fara, D. (2003). Gap principles, penumbral consequence, and infinitely higher-order vagueness. In Beall, J., editor, *Liars and Heaps: New Essays on Paradox*, pages 195–221. Oxford University Press.
- Feigenson, L., Dehaene, S., and Spelke, E. (2004). Core systems of number. *Trends in cognitive sciences*, 8(7):307–314.
- Field, H. (2010). This magic moment: Horwich on the boundaries of vague terms. In *Cuts and Clouds: Essays on the Nature and Logic of Vagueness*. Oxford University Press.
- Fine, K. (1975). Vagueness, truth, and logic. *Synthese*, 30:265–300.
- Frege, G. (1879). *Begriffsschrift*. Halle.
- Fulst, S. (2011). Vagueness and scales. In *Vagueness and Language Use*, pages 25–50. Palgrave Macmillan.
- Gaifman, H. (2010). Vagueness, tolerance and contextual logic. *Synthese*, 174(1):5–46.
- Gómez-Torrente, M. (2011). The sorites, linguistic preconceptions and the dual picture of vagueness. In Dietz, R. and Moruzzi, S., editors, *Cuts and Clouds*, pages 228–253. Oxford University Press.
- Greenough, P. (2003). Vagueness: a minimal theory. *Mind*, 112(446):235–281.
- Hobbs, J. (2000). Half orders of magnitude. In *KR-2000 Workshop on Semantic Approximation, Granularity, and Vagueness*. Citeseer.
- Hyde, D. (1997). From heaps and gaps to heaps of gluts. *Mind*, 106(424):641–660.
- Kamp, H. (1981). The paradox of the heap. *Aspects of philosophical logic*, pages 225–277.
- Kennedy, C. (2011). Vagueness and comparison. In Egré, P. and Klinedinst, N., editors, *Vagueness and Language Use*. Palgrave Macmillan.
- Kirkup, L. and Frenkel, B. (2006). *An introduction to uncertainty in measurement*. Cambridge University Press.
- Lassiter, D. (2011). Vagueness as probabilistic linguistic knowledge. In *Vagueness in Communication*, pages 127–150. Springer.
- McGee, V. and McLaughlin, B. (1995). Distinctions without a difference. *The Southern Journal of Philosophy*, 33(S1):203–251.
- Nouwen, R. (2011). Degree modifiers and monotonicity. In Egré, P. and Klinedinst, N., editors, *Vagueness and Language Use*, pages 146–164. Palgrave Macmillan.
- Pagin, P. (2011). Vagueness and domain restriction. In Egré, P. and Klinedinst, N., editors, *Vagueness and Language Use*. Palgrave Macmillan.

- Raffman, D. (1994). Vagueness without paradox. *Philosophical Review*, 103 (1):41–74.
- Rescher, N. (1969). *Many Valued Logic*. McGraw-Hill New York.
- Ripley, D. (2013). Sorting out the sorites. In Tanaka, K., Berto, F., and Mares, E., editors, *Paraconsistency: Logic and Applications*, pages 329–3. Springer.
- Ripley, D. (201x). Paradoxes and failures of cut. *Australasian Journal of Philosophy*. Forthcoming.
- Russell, B. (1923). Vagueness. *The Australasian Journal of Psychology and Philosophy*, 1(2):84–92.
- Sauerland, U. and Stateva, P. (2011). Two types of vagueness. In Egré, P. and Klinedinst, N., editors, *Vagueness and Language Use*, pages 121–145. Palgrave Macmillan.
- Schiffer, S. (2003). *The things we mean*. Clarendon Oxford.
- Shapiro, S. (2006). *Vagueness in context*. Oxford University Press.
- Smith, N. J. J. (2008). *Vagueness and Degrees of Truth*. Oxford University Press, Oxford.
- Sorensen, R. (2001). *Vagueness and contradiction*. Oxford University Press.
- van Rooij, R. (2012). Vagueness, tolerance, and non-transitive entailment. In Cintrula, P., Fermüller, C., Godo, L., and Hájek, P., editors, *Understanding Vagueness: Logical, Philosophical and Linguistic Perspectives*, pages 205–222.
- Weber, Z. (2010). A paraconsistent model of vagueness. *Mind*, 119(476):1025–1045.
- Weber, Z. and Colyvan, M. (2011). A topological sorites. *The Journal of Philosophy*, 107(6):311–325.
- Williamson, T. (1992). Vagueness and ignorance. *Proceedings of the Aristotelian Society*, 66:145–162.
- Williamson, T. (1994). *Vagueness*. Routledge, London.
- Wright, C. (1976). Language mastery and the sorites paradox. In Evans, G. and McDowell, J., editors, *Truth and Meaning*. Oxford.
- Wright, C. (1992). Is higher order vagueness coherent? *Analysis*, 52(3):129–139.
- Wright, C. (1994). The epistemic conception of vagueness. In Horgan, T., editor, *Vagueness - Supplement of the Southern Journal of Philosophy*. Oxford.
- Wright, C. (1995). The epistemic conception of vagueness. *The Southern journal of philosophy*, 33(S1):133–160.

Zardini, E. (2008). A model of tolerance. *Studia Logica*, 90(3):337–368.

Zardini, E. (2013). Higher-order sorites paradox. *Journal of philosophical logic*, 42(1):25–48.