Tolerance and the Distributed Sorites

Zach Barnett, Brown University, <zachary_barnett@brown.edu>

Consider the family of principles that includes these:

(1) If x grains of sand can be arranged to form a heap, then $x - 1$ grains can also be arranged to form a heap.

(2) If a person with x hairs on his head is not bald, then a person with $x - 1$ hairs is not bald.

(3) If a location x millimeters away from here is within walking distance, then a location $x + 1$ millimeters away is also within walking distance.

These are instances of what are known as *tolerance principles* for vague predicates. If true, they imply that the correct application of such predicates always *tolerates* tiny changes in the objects to which they are applied. Tolerance principles enjoy considerable intuitive appeal: It is strange to think that a heap could be destroyed by removing a single grain, that a person could become bald by losing a single hair, or that a single millimeter could make the difference with respect to whether some destination is within walking distance. At the same time, accepting such principles gives rise to a familiar problem known as the *sorites paradox*: Repeated application of tolerance principles generates absurd outcomes. By repeatedly applying (1), (2), or (3), it seems clear that we will be able to derive – from quite modest premises – that one grain of sand can be arranged in a heap, that a hairless person is not bald, or that the Statue of Liberty is within walking distance of the Gulf of Mexico.

For many, these results constitute decisive reason to think that tolerance principles cannot be universally true.[1] But others are less willing to abandon tolerance principles, sometimes because it is seen as implausible that perceptually indistinguishable objects could be classified differently,[2] or because tolerance is believed to be a defining feature of vagueness.[3] Indeed, some have proposed revising classical logic so that tolerance principles can be held true without inconsistency or absurdity.[4]

My focus here is on a simple such approach, which, in effect, affirms tolerance principles as true but bans their repeated use, thereby rejecting the reasoning (rather than the premises) of the sorites argument. As we will see, this proposal faces a serious obstacle, illustrated by a modified version of the sorites paradox.

---

[1] Fine (1975), Williamson (1994), Sorensen (2001).

[2] Weiss (1976), Kamp (1981), Raffman (1994).

[3] Dummett (1975), Kamp (1981), Fara (2000), Eklund (2005), Gaifman (2010).

[4] Zardini (2008), Weber (2010), Cobreros et al (2012).

1 The 'No Iteration' View

Suppose that we are firm in our commitment to the truth of tolerance principles, perhaps on grounds similar to those offered by Weiss:

> Is it definitely true that any man whose height differs from that of a short man by less than one fifth of an inch is short? I claim that if we assume the men are alike in other relevant ways (e.g., approximately the same weight), then the statement must be true. The reason is that the attribution of shortness to a man reflects our unaided visual perception of the man. We perceive two men who differ in height by less than one fifth of an inch to be the same height. So we must say both men are short or both are not short. (1976, p. 255)

How, then, are we to avoid the absurdity that such a commitment seems to generate? A straightforward proposal is simply to disallow too many appeals to tolerance within an argument. With such a restriction in place, the alleged absurdity can, apparently, be avoided.

Several people advocate this general kind of response to the sorites paradox.[5] For instance, here is Ziff:

> A man with only one penny is poor. That's true. And giving a poor man a penny leaves him poor: if he was poor before I gave him the penny he's poor after I gave him the penny. That's true too. Both of those statements are obviously true. Nonetheless… if you repeat this argument over and over again you'll get into trouble… The moral of the fallacy is plain: it's a perfectly good inference to make if you don't make it too often. (1974, p. 530)

In broad outline, Gaifman puts forward a similar view:[6]

> What is wrong in the Sorites reasoning are not the Sorites conditionals but the way in which they have been put together to produce the contradiction. The unsophisticated intuition that bans the stringing of "too many" Sorites conditionals is essentially correct… Non-philosophers are amused but undismayed by the paradox; they know that each conditional is valid but stringing too many of them is illicit. (2010, p. 11)

We can call this kind of view 'No Iteration,' since it involves, in effect, banning the repeated use of tolerance principles (while still affirming their truth in general). So, for example, recall the tolerance principle pertaining to heaps:

(1) If x grains of sand can be arranged to form a heap, then $x - 1$ grains can also be arranged to form a heap.

The idea is that this principle – while true for any x – cannot legally be used too many times within an argument. How many times is "too

---

[5] Ziff (1974), Weiss (1976), Rosenberg (1978, see pp. 42-43), Gaifman (2010). For a critical review of this approach, see Sorensen (1988 pp. 219-224).

[6] Gaifman uses "Sorites conditional" in place of "tolerance principle."

many"? To this, a variety of answers can be given. To make things vivid, let us suppose that it is illicit to appeal to tolerance more than a single time.[7]

But, here, we must make an important clarification. For, at present, the view described is ambiguous: On one strict reading, it is illicit to appeal to tolerance twice, *even if it is being applied to two distinct objects*; on a more liberal reading, it is illicit to appeal to tolerance twice *only if it is being applied to one and the same object on both occasions*.

It can be shown, I think, that the strict view is unmotivated. Suppose that we encounter two heaps, side by side. And suppose we then take a grain from each. Via tolerance, we can safely infer that the first object is still a heap. So at least one of them is still a heap. It would be strange if, suddenly, we were not permitted to make a parallel inference involving the second object. After all, we could have made these inferences in reverse order.

The driving thought behind No Iteration seemed to be that repeated application of tolerance *to a single object* was problematic – since, in effect, such a procedure would result in a very large change to the object (and no one thinks that tolerance holds for very large changes). But even if tolerance can be restricted so as to block the repeated application to single objects, there seems to be no clear grounds for rejecting its application to many *distinct objects* – so long as tolerance is not applied more than once in each case.

So, the more plausible version of No Iteration, it seems, is the liberal version: It is illicit to apply tolerance *to a single object* multiple times in succession. Henceforth, we will interpret No Iteration this way. As we will see, this version of the view is susceptible to a different sort of problem.

2 The Distributed Sorites

According to No Iteration, tolerance principles such as (1), (2), and (3) are universally true, but it is not legitimate to apply them multiple times to one and the same object. For *reductio*, we assume that this view is correct.

> **Distributed Sorites**: There is a 100,000-step staircase. The bottom step, Step 1, has one grain of sand. Step 2 has two grains. In general, Step n contains n grains (arranged in a heap where possible). The steps toward the bottom obviously do not contain heaps. The steps toward the top obviously do. With respect to some intermediate steps, it is hard to say.
>
> Now, we tinker with the setup: Remove one grain of sand from each step (except for the bottom one), and then add all 99,999 of the grains taken to

---

[7] This supposition is purely presentational.

Step 1, and arrange them in a heap.[8]

BEFORE

AFTER
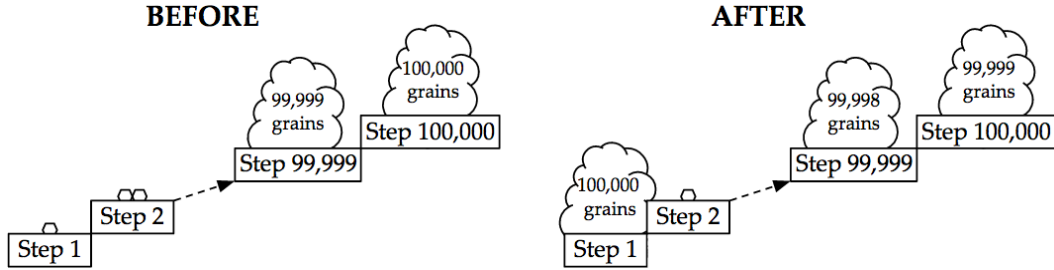
So, what is the problem? By tinkering, we have created a new heap without destroying any. But this is odd, for the new configuration does not seem relevantly different from the original. Let's see this in more detail.

Via tolerance, we can infer that no heaps were destroyed during the tinkering process. After all, we never removed more than a single grain of sand from any individual heap. On the other hand, we clearly created a new heap: Step 1 initially had only one grain of sand, but, after tinkering, it had 100,000 grains. So, we created a new heap on the bottom step, and we did not destroy any heaps that were there before. From this conjunction, it seems to follow that *there are more heaps now than there were before*.

But this is absurd. The new configuration of sand is – in all relevant respects – the same as the original configuration. We began with 100,000 distinct heap candidates, ranging in size from 1 to 100,000. Sure enough, this is precisely what we are left with after moving the grains as prescribed. In effect, we have simply *relocated* each object: The largest object, which was at the top, is now at the bottom; each other object moved up a step, as if on an escalator. In light of the clear correspondence between the new configuration and the original, I take it to be clear that we ought not say that either configuration has more heaps than the other. But this seems to be precisely what the proponent of No Iteration must say.[9]

3 Discussion
Just how problematic is this result? A defender of No Iteration might simply bite the bullet, explaining away any discomfort elicited by the distributed sorites as follows.

---

[8] In debates about ethics, examples with a similar structure have been used. See Parfit (2003, p. 383, fn. 16), Temkin (2012, pp. 440-445), Voorhoeve (2014, pp. 82-84), and Barnett (forthcoming).

[9] It should be noted that Gaifman's (2010) contextualist view is safe from this problem. On his view, a vague predicate "loses its tolerance" in contexts like that posed by the staircase (see his p. 17). The challenge presented is most pressing for views that regard tolerance principles as true without qualification.

Yes, it is true that my view implies that there will be one more heap after the tinkering than there were before it. This fact obtains in virtue of the sort of procedure that was used (i.e. taking a single grain from many distinct heaps and then combining the grains taken to form a new one). Anyone that affirms tolerance will have to admit that the implementation of this sort of procedure will tend to result in the creation of new heaps without ever resulting in the destruction of any.

While it is interesting that the new configuration closely resembles the original configuration in important ways, there is not, strictly speaking, any contradiction – or even any patent absurdity – associated with saying that the former contains one more heap than the latter. Of course, if one could iterate this argument many times over, then it would be possible to derive quite absurd things (e.g. that there are more than 100,000 heaps on the staircase). But, luckily, my view does not allow this.

There is much to agree with here. For one thing, it is certainly true that the proponent of this kind of view can block the extra-absurd conclusions that would follow if the argument given could be iterated indefinitely.

The question, though, is whether the conclusion generated via just one instantiation of the distributed sorites is, on its own, troublesome enough to put substantial pressure on the view. Here, it is worth further emphasizing why this conclusion (i.e. the conclusion that the new configuration contains more heaps than the original configuration) might be aptly regarded as troublesome.

Suppose that we start with two heap candidates, Lefty and Righty, where Lefty contains fifty more grains of sand than Righty. Then, suppose we move fifty grains from Lefty to Righty. No matter what we wish to say about this case, I take it to be clear that *we have as many heaps after this adjustment as we had before it*: If there were none before, there should still be none after; if there was exactly one, there should still be one; if there were two, there should still be two; and, if there was no fact of the matter, then there should still be no fact of the matter in the very same way.

Furthermore, I take it to be clear that these claims are not unique to the imagined situation involving Lefty and Righty: Presumably, these principles can be generalized and extended to cover cases that involve more than two objects. But, crucially, this is precisely what the No Iteration proponent is forced to deny. For, in the distributed sorites, the new configuration resembles the original configuration in just the same way that the Lefty/Righty situation post-adjustment resembles the Lefty/Righty situation pre-adjustment. In such cases, affirming that the total number of heaps does not change is almost irresistible. Denying this seems to me a tough road.

But suppose that we are steadfast in our commitment to No Iteration, in spite of these unseemly consequences. To conclude, it is worth making

a safer and more agreeable point. Before coming to this issue, we might have thought that whether an object is a heap (or whether a person is bald…) should depend only upon features intrinsic to that object (such as how many grains of sand it has, its shape, etc.). Perhaps this is not something that we would hold onto at all costs, but it is, I think, a desideratum. What the distributed sorites shows, at least, is that No Iteration turns out to be incompatible with this desideratum. According to No Iteration, it must be that whether an object is a heap will sometimes depend upon factors extrinsic to that object, such as whether its components came from other heaps. And if this right, then we are owed a story about what these factors are, and why they play this role.[10]

References

Barnett, Z (forthcoming). "No Free Lunch: The Significance of Tiny Contributions," *Analysis*.

Cobreros, P., P. Égré, D. Ripley, and R. van Rooij (2012). "Tolerant, Classical, Strict," *The Journal of Philosophical Logic* 41 (2): pp. 347–385.

Dummett, Michael (1975). "Wang's Paradox," *Synthese* 30 (3/4): pp. 201–232.

Eklund, Matti (2005). "What Vagueness Consists in," *Philosophical Studies* 125: pp. 27–60.

Fara, Delia Graff (2000). "Shifting Sands: An Interest-Relative Theory of Vagueness," *Philosophical Topics* 28: pp. 45-81. Originally published under the name "Delia Graff".

Fine, Kit (1975). "Vagueness, Truth, and Logic," *Synthese* 30 (3/4): pp. 265–300.

Gaifman, Haim (2010). Vagueness, Tolerance and Contextual Logic," *Synthese* 174 (1): pp. 5-46.

Kamp, Hans (1981). "The Paradox of the Heap," in Uwe Münnich (ed.) *Aspects of Philosophical Logic*, Dordrecht: pp. 225–277.

Parfit, Derek (2003). "Justifiability to Each Person," *Ratio* 16: pp. 368–90.

Raffman, Diana (1994). "Vagueness Without Paradox," *Philosophical Re-*

---

view 103 (1): pp. 41–74.

Rosenberg, Jay F. (1978). *The Practice of Philosophy*. Prentice-Hall.

Sorensen, Roy (1988). *Blindspots*. Oxford University Press.

—(2001). *Vagueness and Contradiction*. Oxford University Press.

Temkin, L. 2012. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford University Press.

Voorhoeve, Alex (2014). "How Should We Aggregate Competing Claims?" *Ethics* 125: pp. 64–87.

Weber, Zach (2010). "A Paraconsistent Model of Vagueness," *Mind* 119 (476): pp. 1025–1045.

Weiss, Steven (1976). "The Sorites Fallacy: What Difference Does a Peanut Make?" *Synthese* 33 (1): pp. 253–272.

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

Zardini, Elia (2008). "A Model of Tolerance," *Studia Logica* 90 (3): pp. 337–368.

Ziff, Paul (1974). "The Number of English Sentences," *Foundations of Language* 11 (4): pp. 519–532.