

Beat the Heap!

with *Languages of Art*

Tom Royall 1st August 2015

<http://mathhelpforum.com/math-philosophy/196642-sorites-only-joking-4.html#post744454>

A Tell me, do you think that a single grain of wheat is a heap?

B No. Absolutely not.

A And do you agree that adding a single grain could never turn a non-heap into a heap?

B Kind of – I can agree that no reasonably-sized public opinion poll on the correctness of applying the label *heap* to a range of grain-quantities would show a single additional grain turn an undisputed non-heap into an undisputed heap. Granted, any such poll would have a smallest alleged heap, just a grain larger than some undisputed non-heap. And somewhere along the line – at 274 grains, if you press me now – I shall change my vote. So for me, as for any particular voter, a single grain will appear to make all the difference. Notice, though, that this impression may be misleading. Just as my participation in a political election doesn't necessarily imply approval of the candidate I vote for, so the smallest number of grains that I will vote a heap in a survey doesn't necessarily indicate some sharp personal threshold. Language use, like politics, may be tactical. Nonetheless, however sharp the apparent individual thresholds, their distribution will look increasingly smooth and gradual the larger the poll. I hope this watering-down of your second proposal doesn't miss its point – preserves, rather, its sense of relativism (or gradualism).

A I suppose so. It seems to be one way of blurring the line, which is what my second question seems to demand. What interests me, though, is whether your tampering, in this way, with the second article of faith (the relativist/gradualist one) has compromised your allegiance to the first (the absolutist one). Can you have meant this more complex, relativised sense of *heap*

(that of “popularly considered a heap”), in your first answer? Or did the blurring (the distribution of votes) reach, potentially, all the way back to a single grain? Perhaps, in that case, we must clarify your first answer as “No, absolutely not, *while* I am temporarily inspired with an absolutist intuition of *heap*... whereas, on further reflection, occasioned for example by your next question, I am also happy with a notion of *heap* that recognises a single grain as the very first, albeit smallest, exemplar”. Then, I must admit that the game may be over, since you possibly feel no obligation to square one intuition with the other. I must prepare to be teased that my fascination with the puzzle depends on a curious (but not *that* curious) discomfort with a simple case of ambiguity between alternative senses of a word.¹

I will try, nonetheless, to restart the puzzle, challenging you to confirm or deny this more nuanced and differential reading of *heap*. The puzzle is notoriously robust² in response to such reformulation as you propose, and, “like a virus, will tend to evolve a resistant strain”³. If there is, after all, any important connection between the two intuitions, our time may not be entirely wasted testing them against the next strain. So, tell me, do you consider it a voting matter whether or not a single grain is a heap?

B No. Absolutely not. I agree that the puzzle has effectively restarted, under the new terms of my qualified assent to your second question. You can ask, as you just have, a suitably clarified version of the first, absolutist question, and expect at least a good proportion of interlocutors to still pledge allegiance. Which I do. However, I *did* have in mind my complex, psephological notion of *heap* in my first answer.

I think that we, as people who know how to speak our language meaningfully, will always decline to take seriously any description of a single grain as a heap. We will take such a statement either ironically (probably, as some special kind of paraphrase for a contrary statement) or as merely incompetent – not to be counted an actual use of the word *heap*. The only way that we, our absolutist intuition pumped, could accept that a token of the word *heap* had been made to point at a single grain is by way of reasoning that it denoted in a different language: that the pointing had taken place in its own, perhaps novel system of reference, independent of the myriad pointings that, presumably, comprise usage of the word in English. I propose that we can reflect this absolutist aspect of our semantic linguistic competence in our design of an opinion poll. Specifically, we should choose a statistical framework in which the probability of anyone ever seriously declaring a single grain a heap is assumed to be zero.

I admit that this approach will seem eccentric - lacking in decent democratic scruples, and also in basic etiquette regarding statistical inference. Not that opinion polls (and other surveys) don't routinely, of practical necessity, focus on a finite portion of some potentially unlimited scale of responses, and effectively rule out of consideration more far-flung parts. But they do try not to seem to draw conclusions – whether as evidence or inferences – about those excluded areas. Suppose we approached the present polling project in the normal way, and, in projecting from polls to the distribution of opinion in the whole population, we included in our statistical model a bell-curve (albeit not strictly continuous) with a chopped left tail, in which the probability of recording *heap* votes left of the chop was set at zero.

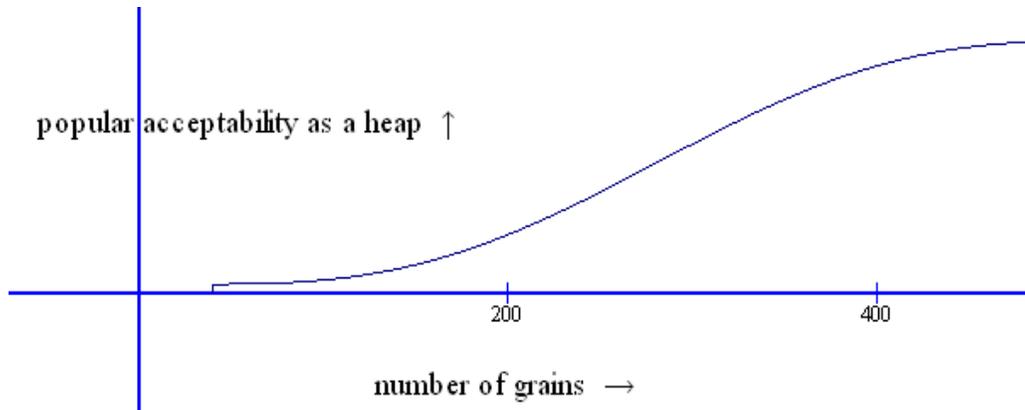


Figure 1.

This could only be (surely?) because we assumed that any probabilities we determined more deliberately (left of the chop) would generally be deemed negligible. But this assumption would be rooted in our experiences with surveys that weren't about the very boundary of the response pattern, and we shouldn't be too surprised to find it contested in the present context. And then we should probably not be concerned to dignify the notion of negligibility, but happy to recast it as some tiny positive range in degree of probability.

But, at the risk, as I say, of being unorthodox, I shall propose what I think are good reasons to set a strict left hand limit on any assumed distribution's support. We shall then be able to justify the very different recasting of 'negligible' as 'absolutely ruled out'. And so – no. I think we can, in principle, rule your question, whether or not a single grain is a heap, a non-voting matter.

A Good. And now, as you probably expect, I shall see whether you are still prepared to assent, with or without qualification, to a suitably updated version of the second, relativist/gradualist question. Do you, now, agree that adding a single grain could never turn a non-voting-matter of *heap*-hood into a voting one?

B Actually, no – I can accept such a step, provided it takes us from a zero probability of *heap*-votes only onto the foot of a very smooth and gradual increase in probability, and provided also that we have a good reason for the location of the step. If I can provide a reason for choosing – say – 12 grains, in preference over its near neighbours, for the honour of marking valid votes from void, then I hope you can agree that the honour itself isn't a problem.

A I hope so too. Please show me the reason.

B We conduct another poll - about an opposing category. As a first approximation, we call it something like *indisputably not a heap*. Or, more pithily, a *pittance* of grains. Any sample for this poll provides, I think, a good reason to choose some particular number as the smallest vote-

able (i.e. disputable) *heap*. If the largest number yet voted an *indisputable non-heap* or *pittance* is 12, then we can set the first (left-most) positive probability of a *heap* vote for 13 grains. And *vice versa* – if the smallest currently alleged *heap* is 38, the last (right-most) positive probability of an *indisputable-non-heap* or *pittance* vote is then to be set at 37 grains. So, just as *indisputable-non-heap* paraphrases *pittance*, a *heap* is an *indisputable-non-pittance*, and likewise puts a limit on the potential reach of the opposing category (figure 2).

You'll have noticed that these polls which are about some categories are, in effect, polls about themselves, and each other. This is, I submit, no problem, if we are ready to refine the approximation in certain contexts. Let's admit that the range of numbers from which the opposing category is guaranteed disjoint is the set of numbers so far alleged to be heaps. Thus, strictly speaking, not *heap* but *alleged-heap* is what deserves paraphrasing as *indisputable-non-pittance*. By the same token, not *pittance* but *alleged-pittance* is what deserves paraphrasing as *indisputable-non-heap*. Also, we should paraphrase as, respectively, *indisputable-non-alleged-pittance* and *indisputable-non-alleged-heap*.

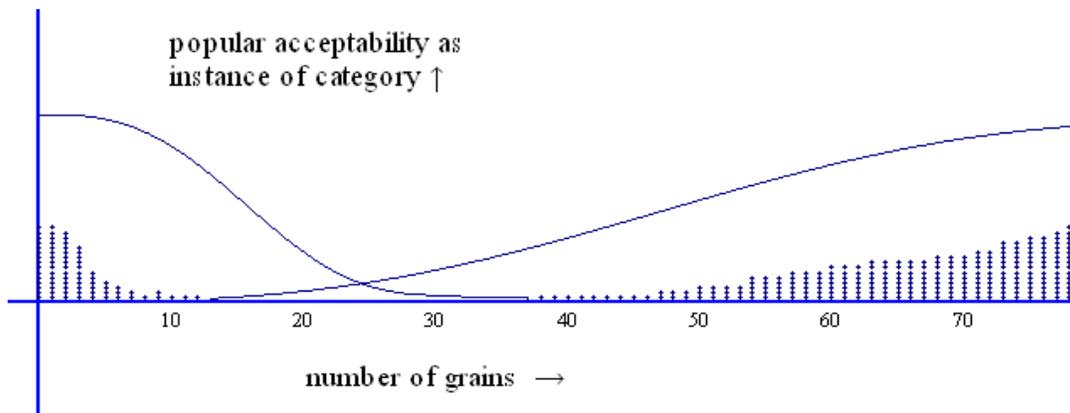


Figure 2.

This leaves *heap* and *pittance* without paraphrase, and I want to reconnect them into the paraphrase loop. The result is more artificial and theoretically loaded than the initial approximation, but I offer it as a method for getting friction on a slippery slope, but yet staying true to the established (though vague) linguistic usage. Firstly, we should recognise *heap* as elliptical, often, for *hereby-voted-a-heap*, and as such not worth tying down to a subjectively neutral extension. And, on the other hand, the poll thus added to is a social construction whose job is to control the extension of *allegedly-a-heap* (one of various less subjective interpretations of *heap*) so that it stays reliably disjoint from *allegedly-a-pittance*.

This is the solution I propose. I think that it satisfies both of the apparently conflicting intuitive commitments that are inspired (in some) by the puzzle. Absolutism is achieved by

treating reference as a matter of social convention, rather than empirical observation. Thus, some votes can be ruled (democratic scruples notwithstanding) out of the question. On the other hand, relativism/gradualism is achieved by treating counting of votes *as* a kind of empirical observation. We are well used to doing so in real opinion polls, of course, and that approach encourages us to trace the further reaches of a category like *heap* as a statistical curve - relative and gradual, if not actually continuous. But then it fails to provide a boundary as such, as the tail extends indefinitely. Whereas, by bending the psephological approach to satisfy absolutism, as described, we get a rather realistic picture of our ordinary experience of categories which are, as often happens, perceived to be intrinsically disjoint. Use of linguistic predicates in this kind of circumstance is explicable as participation in an informal, and very open and accessible – but of course democratically highly dubious – polling system. And vagueness, or semantic tolerance, emerges (as common-sense might lead us to expect that it would) as the lubrication by which conventions can also be customs, which evolve, rather than having to be pre-convened.

The proposal is derived loosely and simply from a concept that Nelson Goodman introduced, in *Languages of Art*,⁴ in order to characterise the behaviour of alphabets, and their semantic counterparts, in various kinds of symbol system. He didn't there, I must admit, refer to votes, probability distributions or the *heap* puzzle. But what I've tried to elaborate in these terms is, roughly, the *finite differentiation* that he sees operating among the inscribed or uttered *marks* (tokens) and *characters* (types) of alphabets and related systems. He defined the principle as follows:

For every two characters K and K' and every mark m which does not belong to both [which, in an alphabet, is every mark m], determination either that m does not belong to K or that m does not belong to K' is theoretically possible. (ibid. pages 135-6)

Goodman observes that in notational systems this syntactic principle can operate semantically as well, sorting objects into categories or kinds that are finitely differentiated from each other, just as marks are sorted into reliably separate characters. Thus, corresponding to each (referring) character of standard musical notation is a class of sounded tones which comply with that character. The compliance classes are themselves finitely differentiated from each other, and comprise, in effect, a semantic 'alphabet'.

It's notable, though not something that Goodman is concerned to examine, that sorites puzzles arise naturally around the borders of 'characters' belonging to both syntactic and semantic 'alphabets'. Finite differentiation provides the intuition of absolutism, and some correlative but much richer ordering of elements than that of relativism/gradualism. Pitched tones, for example, may be arranged in a continuous order of increasing fundamental frequency, as well as falling into a relatively sparse sequence of musical scale-steps. The two orderings correlate, so that if a tone of 250 cycles per second is determined not to be a middle-C, nor indeed any higher scale-step, because it has been judged to belong to the step below, then all other tones up to that same frequency are likewise determined not to be a middle-C, nor indeed any higher scale-step. In such conditions, we can expect a sorites puzzle to form. And indeed, musical sensibility

seems to be against recognising any sharp threshold for qualification as a particular scale-step, yet it confirms equally the absolutist sentiment too. 250 is absolutely not a middle-C; but then how about 251? And so on.⁵

Another kind of example, perhaps more ubiquitous, is gradual progression from one colour into another. The simplest case would be from black into white, where simple luminosity or brightness provides an appropriate gradual ordering of colour instances. This case also has an interesting syntactic counterpart in the case of binary digital coding, where voltage provides the gradual ordering that leads from one to the other character of the binary alphabet. And both cases suggest the possibility of a close analogy with the *heap* puzzle itself. In all three, a binary alphabet (admittedly only latent in the case of *heap*) is traversed by a richer ordering that starts at or next to a zero point, deep in the interior of one character, and continues to indefinitely high values, deep inside the other.

This analogy suggests the following rough interpretation of finite differentiation. To belong to a character (or, more generally, a category or kind) is to be read as such, and we then imagine being able to count readings like votes. Characters (and kinds) are thus conceived as islands of readings (or votes), and their finite differentiation is the assumed guarantee that they never meet.

Turning back again to some of the cases just mentioned, we might be inspired to tidy them up, to the extent of resolving to collect votes on whole classes of tones, colours or voltages, rather than individual ones: just as I have tended to imply that we poll opinion on whole cardinal classes of grain-collections, rather than individual grain-collections. The fact that what we might take to be individual instances of a tone or colour or voltage or grain-collection are themselves analysable, potentially, as sum-totals or classes of individuals (violin-string-sound-events as classes or summaries of ear-drum-sound-events; colour-chips as classes or summaries of qualitative presentations; electrical pulses as classes or summaries of more localised electrical events; grain-collections as classes or summaries of momentary appearances) may embolden us in the quest to classify or summarise. We do, after all, feel free to construe judgements about an individual (so taken) as being about (implicitly) any or all or some things that could be instances of the same thing taken analytically.

Even tidier, though, is to notice that individuation of the votes, relating to frequencies of vibration or to numbers of grains, presents the same choices and betrays the same motives as individuation of tones or grain-collections instantiating those classes. A particular vote alleging 38 grains to be a heap might correspond to a particular collection of that number, perhaps a peculiarly *heap*-ish exemplar of its class; and a vote for putting the frequency of 255 Hz into *middle-C* might correspond to a particular tone, perhaps from a peculiarly sensitive and daring performance. We might, for some purposes, feel obliged to collect votes on the particular grain-collection, and the particular tone; but if, on the other hand, we are content to refer at once to the whole cardinal class of grain-collections or the whole frequency class of tones, we might well find we can keep our options open as to whether to require that votes be cast by voters, in the normal way, or perhaps instead by exemplars of the class. As Goodman brings out in his theory

of exemplification, the means by which a label can be made to 'apply to' or to 'point at' a referent are far from being exhausted by explicit verbal description, and are often achieved through mere (but contextually apt) display of an instance, which thereby as much points out the label as gets pointed out by it.

So, against any *prima facie* implausibility in the notion of counting votes in order to solve the puzzle we can weigh a remarkable simplification which is available from letting votes be exemplars. We are able to picture, for instance, *middle-C* as a bell-shaped distribution of tone-votes over the frequency continuum, such that those of relatively eccentric pitch are relatively few. Such a picture, or that of a whole scale as an archipelago of island clusters, appears preferable - perhaps more musically plausible - than one in which each step has final boundaries, and no interesting topographical profile. Now, you may be coming from a different musical or theoretical direction, and suspect that I am confusing the distribution (along the frequency continuum) of *middle-C* tones as such (which you do suppose to have a roughly uniform profile) with the distribution of a specific subset of those tones: those that are aesthetically worthy, or are the produce of expert performers. If so, we understand each other well enough. I merely question the theoretical need for the sharply defined scale, and for a category *middle-C* that is distinct from *good-middle-C* or *acceptable-middle-C* (or whatever the title that you prefer for what I call *middle-C*). And this significant theoretical economy comes from my ability to regard a tone as a vote for including its location, on the frequency-continuum, within the scope of *middle-C*.

Or, consider ethical categories such as *person*. We can at least conceive of collecting data on how such a category distributes along a rich ordering such as that provided by, say, number of seconds since conception. We might, even, want to tackle the slippery-slope nature of ethical debates about person-hood, by assessing the possibility of finite differentiation in such category-pairs as *person* vs. *human embryo*. In positing such a "semantic alphabet", with no grounding in a corresponding actual, syntactic alphabet, I depart rather drastically from any explicit recommendations in Goodman's theory. Nonetheless, perhaps ethical disputes could benefit from serious consideration of whether such an alphabet were somehow latent in related discourse. This might involve polling opinion, though not necessarily in the normal and explicit manner, even granted my eccentrically undemocratic statistical framework. Rather, in a society in which legal courts are allowed discretion in passing judgement on particular disputed cases, a record of past judgements might (with appropriate extrapolation) instead provide the distributional picture we require. Here, each case may seem to (in a sense) argue, cast a vote for, including that particular age in a particular category.

From this point of view, the familiar lament that, given a rough correlation of a given rich ordering with membership vs. non-membership in a category, we then "have to draw the line somewhere", thus perfecting the correlation, seems unnecessarily defeatist. We are all the more likely to accept (as the puzzle's second question already encourages us to) the imperfection of this particular correlation the more we understand how it can help maintain a perfect one: that of the same rich ordering with membership in the characters of a semantic 'alphabet' (whose finite differentiation satisfies the rather stringent absolutist demand of the puzzle's first

question). Thus, some individuals judged to be *non-persons* will be older than some that have been classified *persons*. But this room for discretion might, conceivably, be what enables the legal system to achieve finite differentiation of ethical categories, and guarantee that no individual judged an *embryo* will be older than any that is judged a *person*.

A In that case, I must be careful to clarify, at the start of the puzzle, that I am not at all talking about whole classes of grain-collections, but rather any one such collection, which I ask you to imagine augmenting, grain by grain. Do you have a way, now, to deny that any single additional grain will turn this particular collection into a heap, without readmitting the psephological complications that you were trying to dispense with? Or to deny, without readmitting those complications, that an increase of a single Hz in frequency will bring a particular glissando-ing tone into the territory of middle-C, or that ageing by a particular second will turn a particular individual into a person?

B I do have a suggestion, though I doubt that it entirely satisfies your proposed restriction, and doesn't imply, after all, a potential multiplicity of verdicts for each grade in the gradual progression. And I wasn't trying to avoid such multiplicities but pointing out, merely, that they might be neatly provided by exemplifications instead of straight-forward labellings. And you evidently were not previously disturbed by this avowed pluralism, only concerned (quite rightly) about whether it compromised my absolutism. Still, I see that I have, throughout, rather encouraged the assumption that the voting pattern of each voter (or exemplar) should be expected to correlate perfectly: show a single and sudden threshold. I think this assumption is simplifying and harmless, unless and until we try removing the pluralism as you are now apparently inclined.

My answer is to point out that the assumption is, though largely harmless as long as the pluralism holds, unnecessary. Each person may, quite plausibly (but requiring a more expensive survey design), change their vote plenty of times over, so that, for example, 30 grains is their smallest heap, 52 their next smallest, then 57 and 58 are both heaps, then 60, then 70, then 72 and 73, then 76 through to 79, then 82 through 95, then 98 through 110 then 112, and their largest non-heap is 131. I, at least, might well submit these votes, if the survey design permitted. Or I might create some even more complex and smoother increase in the density of my heap votes from left to right. I don't see that a single, authoritative (e.g. legal or scientific) series of judgements shouldn't make a similar pattern. You would be right, of course, to construe it as evidence of an underlying pattern of voting dispositions, and then to restart the puzzle to ensure that absolutism wasn't thereby compromised. And then I would invoke a terminating projection curve, and hence the pluralistic notion of a multiplicity of future (or yet to be counted) verdicts at each step of the gradual progression. But then I would hope your aversion to this approach had subsided.

A I see. Tell me, though, do you need any of this wealth of polling evidence to decide whether a single grain of wheat is a pittance, or, indeed, whether a trillion grains is a heap? And, if not, can you specify the number of grains at which you first begin to require such evidence?

B No, I don't need anything like an actual poll; only a very large but hypothetical one - a notion which I think is very commonly conjured when we refer to "usage". For example, I shall justify my judgement that a single grain is a pittance by explaining that my general impression of past usage of the word, however hazy the impression, and without necessarily recalling any specific instances, makes me confident enough that a single grain would be among the highest-scoring numbers (of grains) in any sample that we did bother to collect, somehow, of actual uses of the word. Competence in using the term brings the confidence to speak, at least in clear cases such as this, on behalf of the whole society of competent users. I'm equally happy for the enquirer to read my affirmation, in the clear case, either as a single new vote or, in the absence of more scientifically grounded evidence, as anecdotal evidence that a very large number of similar votes have been observed. So I set up a poll of a kind, right from the start. And I don't have to address the question where precisely the need would arise further on.

You might ask: at what point (what number of grains) am I no longer served by the lazy, egocentric survey, so that I require something more solid and public? But the question is no more disconcerting than having to locate my own personal threshold. And, incidentally, my answer is no less sympathetic to the relativist/gradualist spirit of the question. It doesn't give an arbitrary cut-off point but one arrived at from examining a rich picture of data. The numerical foreground of the picture - the count of votes - will usually be sparse. But the context of each vote, which is what I know of the voters, however few, and other potential voters, and all of their circumstances and motivations, is still rich enough to make my answer into a rational judgement call.

In the absence of relatively hard or publicly agreed polling evidence, the two answers will be the same. My threshold will be the largest number of grains that I am happy for you to record as getting either my vote or, on merely anecdotal grounds, some large number of similar votes. Being thus ever-ready to downsize but solidify my claim, from anecdotal and general to immediate and personal, I shall want to imply that my threshold bounds a set considerably smaller than the actual extension of the word. But this, also, is no embarrassment. I am simply ready to retreat to a clearer and more common-sense distinction between observed and unobserved, that is, sample and projection.

In an egocentric, bootstrapping poll, the sample is confined to my own votes, and therefore terminates at the same number of grains as my personal threshold. The sample could even peter out rather than distribute uniformly, if I offer a rating of my certainty for each number of grains. But, at my personal threshold number of grains, the rating will reach a zero value. In a larger poll (perhaps, next of all, a pooling of mine with your own bootstrapping poll), the sample is larger than my own contribution, and my threshold probably smaller than the largest number of grains so far voted a pittance.

A Okay. But I notice that your initial, boot-strapping survey, though you call it egotistical, is not terribly assertive, and shrinks, after all, from absolutism. It asserts merely that, in relating *pittance* to a single-grain, any real poll would have votes for far outnumber votes against. This suggests that you are not ruling out a proportion, perhaps tiny, of votes against. So you are

prepared, in principle, to accept that a competent speaker might decline to vote a single grain a pittance. And, presumably, you would be willing to record the vote against as equivalent to one in favour of applying *non-pittance*. Is this right?

B Yes, roughly speaking. I didn't quite admit that a real poll would have to force each respondent to choose between a vote for and a vote against applying some label at each (or some) number of grains. I alluded rather to the possibility of counting uses of the positive label only, and examining their distribution over different numbers of grains. But, by the same token, we might equally well count uses of a corresponding negative label, and then we should expect the two distributions to overlap, with at least roughly the same complementarity that would result from forcing a choice for each respondent – so that one distribution would build as the other subsided. This will have the consequence which, understandably, gives you pause. The projection curve for *non-pittance* will have a profile similar to that of *heap*, but penetrate not only the projection curve of *pittance* but also right through any sample of *pittance*, all the way back to a single grain.

I agree that such a profile deprives *pittance* of absolutism, in the form of a core region reliably separated from *non-pittance*. But who said that absolutism in this form is required for *pittance*? The role for this category, which it plays admirably, is to be the core region of *non-heap* that is reliably separated from *heap*. Or do you, indeed, want to restart the puzzle by asking me if I could countenance a single-grain non-pittance? I doubt that a single grain is far enough from the (fuzzy) border of *pittance* to make assenting to that particular version of the first question sufficiently unthinkable.

A I see. But in that case, I certainly do have a question. Can you be as accepting of the possibility of someone properly calling a trillion grains a non-heap?

B No, should anyone start the puzzle from that end, and (as seems likely) ask about *non-heap* rather than *pittance*, then I shall point out that a trillion grains is a clear enough case of a *heap* to be considered a *mega-heap*, which can stand to *non-heap* exactly as *pittance* does to *heap*. *Heap* does have, potentially, a core region which use of an intensifying label like *mega-heap* can reliably separate from *non-heap*. So I can be absolutist at first (and quasi-relativist-gradualist secondly) just as before. But the latent 'alphabet' of categories that I will try to nurture to a plausible state of maturity now consists of *non-heap* and *mega-heap* (plus middle-ground), instead of *pittance* and *heap* (plus middle-ground).

The presence of *heap* in one alphabet and *non-heap* in the other suggests that the two ought to merge. A composite of the two is indeed viable, and worth outlining. One reason to do so is to depict our recent observations about the relation between *heap* and *non-heap*, and show that it doesn't cause any logical inconsistencies, despite (or because of) its requiring the prefixing of both terms by either *allegedly-* or *hereby-voted-* (or, reporting particular speech acts, *thereby-voted-*), which ordinary discourse can often do without. Another reason is to show an interesting comparison among musical scales and colour scales, and even ethical scales and physical measurement scales.

To this purpose, it will help to view the regions of figure 2 from above, in a kind of Venn diagram. First of all we should imagine each 'island' of votes as (ironically!) a heap of grains on a flat surface. Such an image is encouraged by the statistical method of cluster analysis; we of course add the absolutist assumption that clusters are entirely separate, and that the projected reach of each stops at the present reach of the next. At least while the domain of elements can, as with numbers of grains it can, be restricted to a single dimension, the simplest overhead picture is then formed by smooth 'shorelines', actual and projected, for each island, as in figure 3.

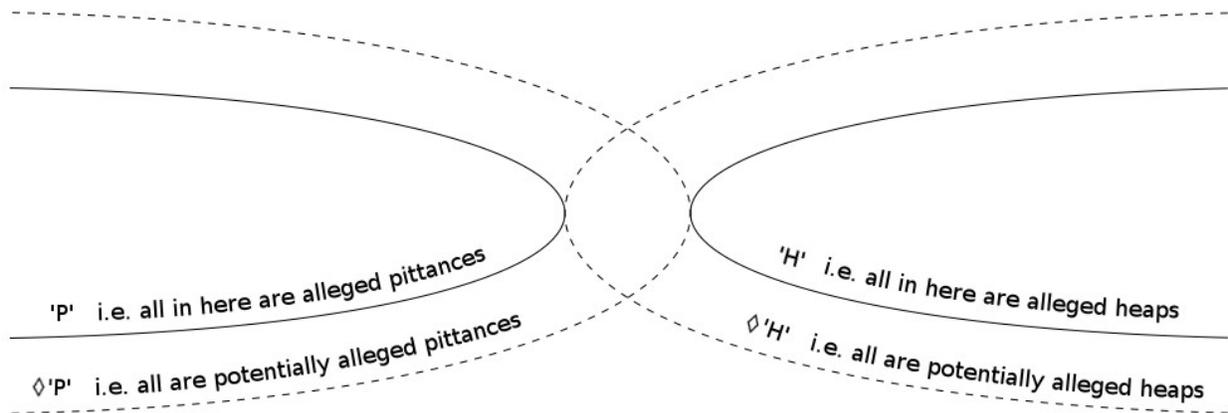


Figure 3.

Into this picture, we can clearly enough incorporate a poll for each of the corresponding negative labels. These can represent usage not only of the negative labels themselves but of their conjunction, *neither-heap-nor-pittance* (bold, in figure 4). Parties in a slippery slope dispute could, I think, find it useful to recognise such a 'middle-ground' category, between two others, that differs from the actual current middle-ground in overlapping with them.

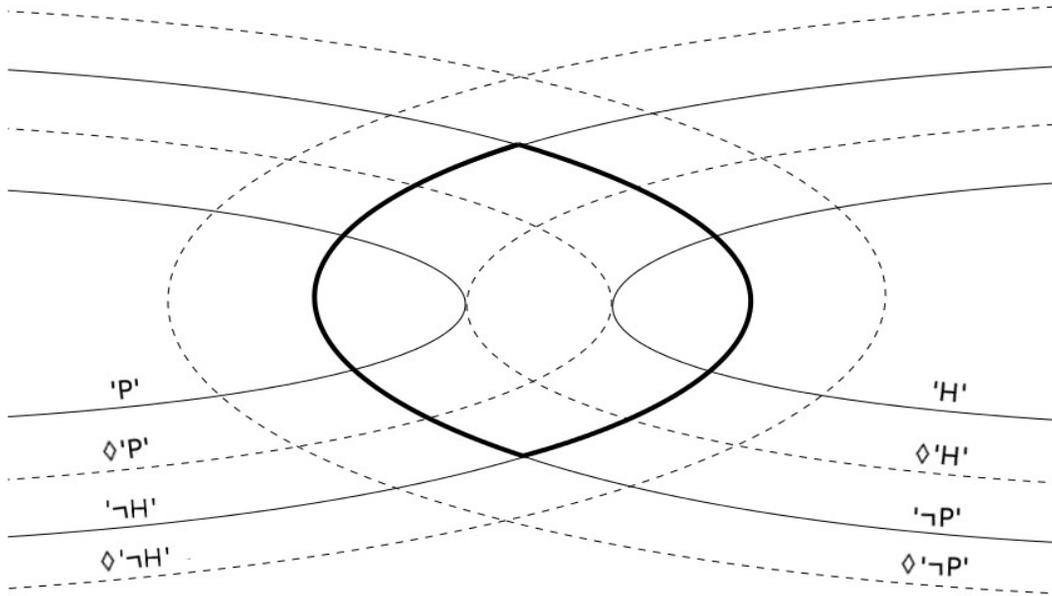


Figure 4.

Having now incorporated the *non-heap* character of the second alphabet, we shall use the *mega-heap* character, too. First without *non-pittance*, so that we have just the composite picture as earlier promised:

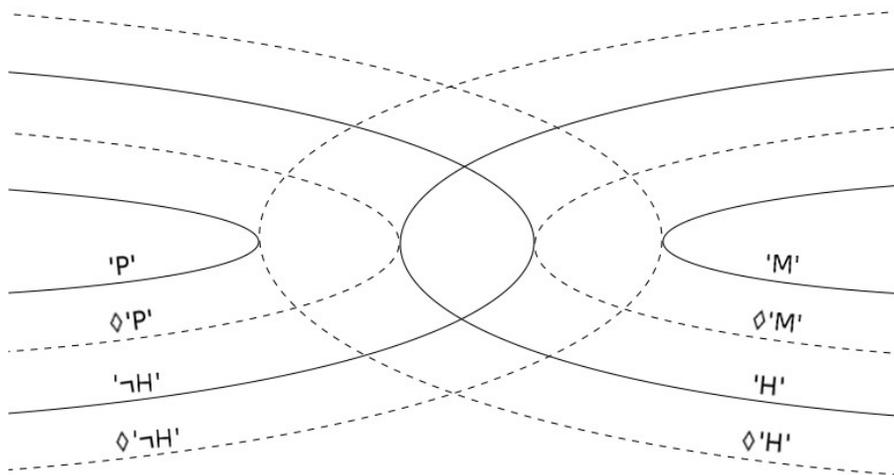


Figure 5.

And now, in figure 6, we include negatives for both *pittance* and *mega-heap*, so that two conjunctive categories form, and begin to indicate the parallel with musical and colour scales.

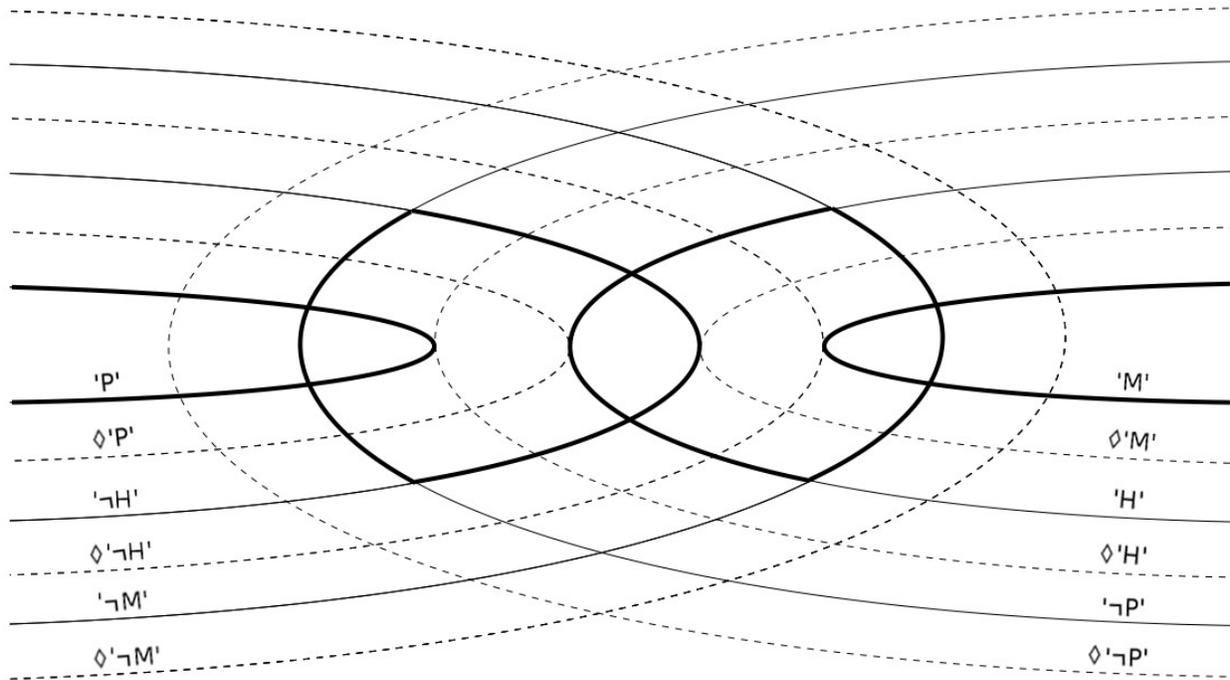


Figure 6.

The parallel with colours is clear enough. The bold regions link by overlapping to make a chain, just as colours might link by overlapping to form a rainbow. Relative to the basic labels *pittance* and *heap* the bold regions are (*allegedly-*) *a-pittance*, *neither-a-pittance-nor-a-heap*, *a-heap-but-not-a-mega-heap*, and *a-mega-heap*. The corresponding colours might be: (*allegedly-*) *orange*, *yellow*, *green*, and *blue*. This, then, is a way of showing usage of *orange* to be finitely differentiated, semantically, from that of both *green* and *blue*, and *orange* and *yellow* both similarly from *blue*. But it also shows that *yellow*, as well as being differentiated from *blue*, is a way of referring to the middle-ground that might well differentiate *orange* from *green*. (And *green* refers to the middle ground between *yellow* and *blue*.) It (*yellow*) is bigger than the actual middle-ground between *allegedly-orange* and *allegedly-green*, and overlaps them both. So it lacks differentiation from either. And, consequently, it cannot be a character in the same alphabet. What we do apparently have, though, in a rainbow perceived with a child's colour vocabulary, is a pair of interlocking semantic alphabets: *red-yellow-blue-violet* and *orange-green-indigo*.

Arguably, a musical scale is one of such a pair; but there are striking differences here.

The other, interlocking scale/alphabet is suppressed or discarded. We don't give names to the 'bad' regions of pitch between scale-steps, and would find it difficult, at least at first, to recognise and compare them. Also, the pitch steps are fairly narrow, compared to the gaps between them.

Likewise in physical measurement to a set precision such as, say, one decimal place for a stated unit. Adjacent values on this scale, for example, 2.3 and 2.4, cannot be regarded as semantically differentiated if they represent, as is often inferred from the expression “correct to the nearest tenth of a unit”, the intervals [2.25, 2.35[and [2.35, 2.45[. However, semantic differentiation is, arguably, required of such a system in order to allow replication of experiments,⁶ and it might be achieved in either of the two ways just compared. The scale could contain two interlocking semantic alphabets analogous to the rainbow just considered. Only one is used, the other discarded, its compliance classes corresponding to syntactic 'malformed' or 'non-character' marks. Then the scale really imposes increments of one fifth of a unit, not one tenth. Or, more like a musical scale, the supposed representation of intervals could be ignored in practice, and objects chosen for measurement that cluster close to the ideal values 2.3 and 2.4. Either way, though, while 'open texture' is inevitable between any value and its negative, differentiation is generally achievable between adjacent semantic 'characters' that are widely enough spaced.⁷

A That may be all very well, but I'm not yet happy that you've squared the intuitive circle. I am pleased that you want to put clear blue water between your smallest possible alleged-heap and zero grains. As I lamented earlier on, many people, while disinclined to have the curve extend indefinitely to the left (supported by negative numbers of grains, perhaps), are happy for it to begin its ascent from a zero probability of alleged *heap*-hood (or, as they might prefer to have it, a zero degree of *heap*-hood, or a zero degree of truth of *heap*-hood) at precisely one grain. In such cases, the game ends badly. I accuse my opponent of simply not seeing the point of the puzzle, and (infuriatingly enough) they usually agree. My sense of relativism/gradualism in these cases remains untroubled. I agree with you that the puzzle's second question dares us to provide a good enough reason for locating the step at some particular number of grains, but this kind of opponent has the undeniably forceful reason that one grain is at least some wheat as opposed to none at all. However, in such cases, we don't even get to the second question, because my absolutist sensibility is too bruised.

Now, I appreciate the trouble you've taken to achieve an adequate absolutism, but it seems to be in vain. Absolutism does extend admirably throughout your current poll of votes alleging *pittance* and thus implying *indisputably not an alleged-heap*. So that, were we to play the 'iterative' form of the game (in which, rather than asking the second question in generalised form I ask instead whether you think that 2 grains is a heap, then 3 grains, and so on) you could answer as confidently the 12th time around as the first. Well, stepping onto a positive probability at 13 grains, I'm sure you won't be surprised if I still harbour relativist/gradualist doubts about the step. But I must say I wonder also whether our absolutist allergy to having the *heap* projection curve's support get too close to zero grains isn't just as triggered at this step. You wanted some clear blue water between the smallest possible heap and zero, and now don't you

need it likewise between the smallest possible heap and your new category? And, if that weren't bad enough, the projected boundaries of the two categories even overlap! So, just how may you claim to be defending absolutism?

B I hope I can explain. First, notice that, although the adjacency certainly needs explaining, the overlapping is no worse. Where the curves overlap, no votes have yet been recorded in any poll, that we know about, for putting these numbers of grains in either category. So this ground, being as yet unclaimed for either territory, is potentially the property of either. Though, only as long as it remains unclaimed for either. We could say that each of the two probabilities represented by the heights of the curves above any particular number of grains in the no-man's land is conditional on the other turning out to be zero. So each of the curves in this current middle zone assumes an eventual retreat of the other one out to the number marking what is, currently, the edge of this zone.

I superimpose the curves because the two retreats are equally plausible, as things stand - which is to say, actual recorded polls being their present size. If relatively eccentric, extremist votes ever succeeded in expanding the territory of either category, so then the range of possible expansion for the other category would retreat. But either direction of travel is theoretically possible. We might assert that a more plausible eventuality is a partial movement both ways, so that both polls advanced, and curves retreated, to somewhere in the middle (figure 7).

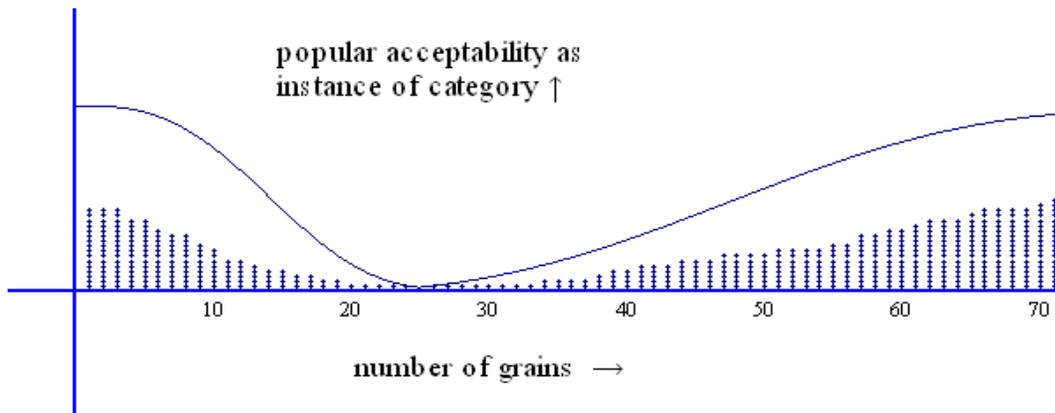


Figure 7.

Which is true. It just fails to indicate the present boundaries of possible future growth for each category.

Perhaps this doesn't prove what I said about overlapping (of the two probability curves) being no worse than adjacency (of either one of these curves to an 'opposite' poll or sample)... but I hope it deodorises the former. As for adjacency itself - this becomes much less alarming if we allow that the probabilities in the middle zone are very low indeed. They make us fairly sure that territorial gains up to most (and quite possibly all) of these points will never actually occur,

during any realistically conceivable lifetime of the polling system. So the middle zone is itself clear blue water (... a margin for error, a buffer, no-man's-land, elbow room, slack in the system) between the polls as they stand. It equips the system to survive an emergency, in which the zone were to either disappear or threaten to. But the system is almost certain never to encounter such an emergency, in its finite life-time. The curves, then, are only a theoretical abstraction to help explain how and why we can depend – for practical purposes – on the perpetual separation of the polls themselves.

A But 'most of', here, sounds very suspicious. Just to be clear, do you think we can be fairly sure that **all** points in the no-man's-land are immune from occupation?

B I suppose not. But, much as a *sorites* sceptic might cheerfully trade in *heap* for a gradual scale of relative *heap*-ness (starting with grade 1 at a single grain), I'm happy to swap my fair certainty here for at least 25 grades of the same. Like her, I shall regard the binary choice as just a first approximation, where it draws the line perhaps pointing towards features (such as central tendency and skew) of a more informative grading, but not bounding a territory in which any absolutism is asserted.

I deny that this makes me, like her, a *sorites* sceptic. On the contrary, I quite sympathise with the suspicion that I might, conceivably, have been talking about a category *fairly sure of...* that, despite being qualified by *fairly*, was actually finitely differentiated from some other category (such as *not at all sure of...*), and so behaved with an absolutist dimension just like *heap*, and hence coloured paradoxical when put to the *sorites* test. But my priority at the moment is to explain this kind of linguistic paradox, and we don't need a new example, in addition to *heap* itself. So I shan't be encouraging that reading. Rather, I'm happy that the 'back-end' of my opposing-polls explanation should be tied up with just the kind of trade-in that would have disappointed both of us at the beginning.

A Okay, I accept that if you step into the water at the edge of an island, thereby relinquishing any absolute guarantee against long-term encroachment by the next island, your relinquishment will be genuine. You aren't about to embrace the same guarantee in a new guise. But I think you underestimate the difficulty of relinquishing at all. You put the difficulty down to an intuitive embarrassment over privileging any particular number with the honour of marking the crucial step. That may be right – partially, and initially. But now your theory, at least, presents a purely logical reason to doubt that the step is feasible.

Let's stand at the edge of the current poll for *pittance*. We are either on land, at 12 grains, or in water at 13. Now let's consider the island immediately behind us, instead of the projection curve projected from the other island, some distance ahead (though the projected curve itself is present - or, if we are still dry, adjacent). The island right under or next to us is made of actual, accepted votes, not the hypothetical and credulity-testing kind that seemed to be in question at the outer reaches of the projection curve, and provoked your dictatorial approach to opinion polling. Now, even a single recorded vote for 12 grains as *pittance* makes it an *alleged-pittance* or *indisputable-non-alleged-heap* and thus absolutely prevents it from being (alleged) a heap,

yes?

B Supposing the vote to be valid, then yes, that's the idea.

A And obviously, since 12 grains is as well and truly immune from being a heap as has been alleged, any such allegation (with respect to the same number) is equally warranted. I feel bound to offer my vote, then, for 12 grains being a *pittance* and therefore an *alleged-pittance* and therefore an *indisputable-non-alleged-heap*. And you will, surely, accept it. But then, in my position of absolute certainty about 12 grains, as a validated voter of 12 grains a *pittance*, which makes it an *indisputable-non-heap*, I surely can't fail to register an opinion I have about 13 grains. If that opinion falls on the *pittance* side, and I don't see why it shouldn't, then 13 becomes allegedly a *pittance* and hence actually an *indisputable-non-heap*, and 14 grains soon after. So your poll then goes out of control, and swallows up the margin for error.

B I agree that we are often worried in this way about margins for error. It's also noticeable that the worry doesn't usually result in a margin actually shrinking. Nonetheless, theoretical as well as practical resolution is available, if we take care to distinguish between *allegedly-a-pittance* and *hereby-voted-a-pittance* as recommended earlier (with respect to *heap*).

“... as well and truly immune... as has been alleged” is where the distinction is needed. It's true, and remarkable, that one or more allegations are the source of immunity. Also that any other allegation of immunity must then be equally valid. But any such would be an allegation that 12 grains is an *alleged-pittance*, and this is not the kind of allegation that gave rise to the immunity. Indeed, one benefit of polling opinion on an issue is that we usually feel justified in not bothering to poll opinion on the count itself. We would demand unanimity on it, so that any current poll of votes affirming the correctness of applying the label *alleged-pittance* to a range of numbers had a uniform profile. (Figure 8.)

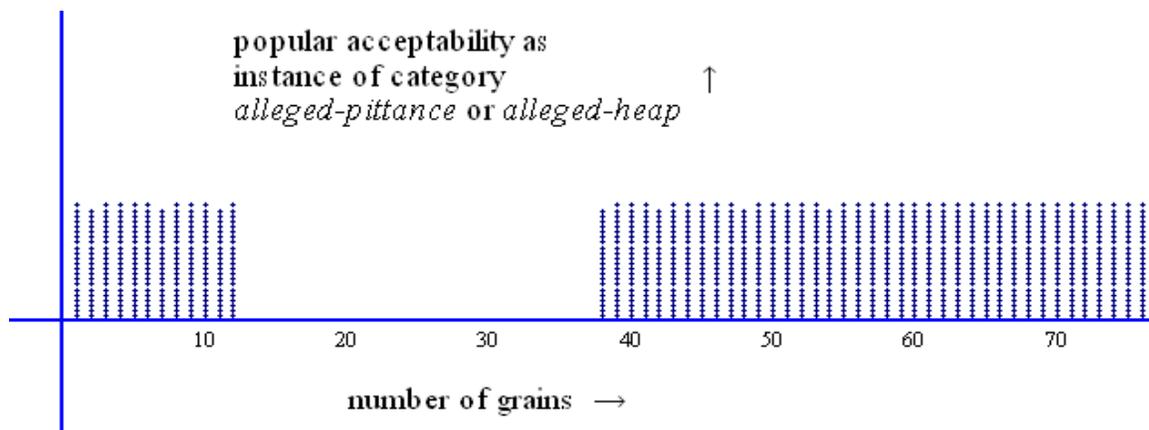


Figure 8

Where we expect such a profile, we expect to benefit from the unanimity of discourse about the polling data, but not from bothering to actually conduct the meta-poll. And I too am not interested in the meta-poll. I think you have mistakenly thought that I am. I admit that it might seem natural to suppose that, if anyone should be interested in the meta-poll, it would be me! But not so. The allegations that interest me are such as those alluded to by the label *allegedly-a-pittance* (uses of *pittance* itself), because it is these which, were they to build up along the current border of *allegedly-a-pittance*, would be likely to spill outwards into the margin as well. Where, as with "... as well and truly immune... as has been alleged", we fail to distinguish the two polls, we might easily imagine that the actual allegations of *pittance* would tend to build up at the border (like allegations of *allegedly-a-pittance*, where the discourse is matter-of-fact), but then tend, consequently, to also spill over (as might similarly-stacked allegations in a matter of opinion).

But we shouldn't expect such a build-up. Let's examine the votes you felt bound to offer, and sure I would accept. I will certainly accept, for the meta-poll if any, your vote for 12 grains being an *alleged-pittance*, and I agree that this makes this number an *indisputable-non-alleged-heap*. However, while both of these would follow from success in voting the same number a *pittance*, the reverse isn't the case. This latter vote is the one I shall have to refuse to accept. Anyway, you will probably withdraw it. Bear in mind, we have just agreed that the current profile of *pittance* votes peters out at 12. A vote for 12 (as *pittance*) is thus highly eccentric, though not unprecedented. You will be hard put to convey such a highly eccentric (even though not unprecedented) opinion as genuine and sincerely held by yourself - now, in the middle of this puzzle - because less pure motivations are much more plausible. So, whether simply embarrassed, or else consciously striving to maintain the benign atmosphere of clarity and matter-of-fact-ness that issues from unanimity on questions of vote-counting, you will probably retreat to a more plausible threshold. The current profile then still peters out, as it did before. The vote-count at the border hasn't increased, and we are no more inclined than before to expect the border to advance.

A I see. So, as well as presuming to rule out any votes that would cause opposing categories to overlap, you will charm, bully and cajole to discourage highly (and perhaps even mildly) eccentric votes, too. But what happens if I decide to stand my ground, and not (which you admitted was only probable) withdraw my eccentric vote?

B Actually, I might then accept it. There was a degree of bluff to my previous stance, and I could now be impressed by your willingness to call. This doesn't mean I have to concede the domino effect you claim. I will, rather, have judged that the middle ground is large enough to take the strain of this, and possibly many similar votes. If only the one, then the poll still tapers to the same point, only with added 'fuzz' or texture. If many, then the new, larger sample will taper out further out. This time, the chances of your persuading me to see it extended, here and now, will be negligible.

A Exactly. And what then, when you have run out of bluff? What if we *both* stand our ground? I don't imagine you can seriously deny the fact (in all of its clarity and matter-of-fact-

ness, as you put it, although associating the two seems questionable) of my at least offering the vote. Shouldn't you at least register it in a poll of attempted votes, of which the accepted and validated ones form a sub-set?

B Yes, if I saw reference only in an atomistic way, built, bottom-up, from individual votes, the token instances of reference. Then I should be hard put to deny that votes properly voided from some particular poll or sample must deserve counting nonetheless. I might have the clearest political or psephological grounds for excluding them from some such particular poll or sample. But if the poll or survey is designed to throw light on real usage of a term in natural conditions, and if I assume that such usage boils down to individual acts of reference, of pointing of the term at one or more objects, then you would think that I can hardly deny that your attempted vote must, as you suggest, qualify for inclusion in some differently motivated sample. And then the prospect emerges of a chain of gradually more inclusive polls or samples. So I seem to need to justify an ostrich-like attitude of wilful ignorance towards the empirical data you offer.

Now, I do see reference in this atomistic and bottom-up way, but only in one side of a circular, holistic story, in the other side of which the atoms are themselves inferred from the top-down influence of type reference. This story helps me to see that your attempted votes are not *necessarily* worth counting, in any poll at all, and that my attitude is one of useful focus rather than wilful ignorance. I think that such a story guides ordinary language use, in that we continually hint at either or both kinds (upward and downward) of explanation for our utterances. We validate a predication by noting that the described object satisfies the adjective as the latter already stands, as an established enterprise (top-down). And we also audit the enterprise, assess the scope of the adjective, by counting or pointing out individual descriptions as given entities, as though they stand in their own right (bottom-up). I must concede that we don't, ordinarily, follow through on these hints, perhaps because the holism envisioned is circular and complex. But I suggest that they do bear elaborating, perhaps most clearly in the case of notational reference.

WIP ;)

Notes

1. E.g. Richard Dawkins' tweet: "The sabbath was made for man, not man for the sabbath.' Sorites is paradox only if you forget that words are our servants not our masters." Twitter 25/09/14, <https://twitter.com/RichardDawkins/status/503810329382354944>. See also "The Tyranny of the Discontinuous Mind", in *The New Statesman*, Christmas 2011, <https://richarddawkins.net/2013/01/the-tyranny-of-the-discontinuous-mind-christmas-2011/>

2. Terence Horgan, “Robust Vagueness and the Forced-March Sorites Paradox”, *Philosophical Perspectives*, 8, *Logic and Language*, 1994,
<http://thorgan.faculty.arizona.edu/sites/thorgan.faculty.arizona.edu/files/Robust%20Vagueness%20and%20the%20Forced-March%20Sorites%20Paradox.pdf>
3. R. M. Sainsbury, “Concepts without Boundaries”, 1990, in *Vagueness: A Reader*, Rosanna Keefe and Peter Smith (Eds.), MIT Press 1999, page 260.
4. Nelson Goodman, *Languages of Art: An Approach to a Theory of Symbols*, Hackett 1976.
5. Musical sensibility perhaps also urges a reformulation in terms of relative rather the absolute pitch. By my reckoning, this would not greatly affect the shape of the resulting sorites dialogue.
6. Catherine Z Elgin, “Relocating Aesthetics”, in *Between the Absolute and the Arbitrary*, Cornell University Press, 1997, page 75.
7. Israel Scheffler, *Beyond the Letter*, Routledge & Kegan Paul 1979, pages 60 – 64.