

Lessons for vagueness from scrambled sorites

Abstract: Vagueness demands many boundaries. Each is permissible, in that a thinker may without error use it to distinguish objects, though none is mandatory. This is revealed by a thought experiment – scrambled sorites – in which objects from a sorites series are presented in a random order, and subjects are required to make their judgments without access to any previous objects or their judgments concerning them.

§1. Scrambled sorites

A *sorites series of objects* for a pair of predicates, P1 and P2, is a collection of objects meeting the following conditions:

1. At least one of them is P1: for some i , $\#i$ is P1.
2. for some i , $\#i$ is P2
3. for no i is it the case that $\#i$ is P1 and $\#i$ is P2
4. Applicability of P1 and P2 depends upon some property F that comes in degrees (how red, how tall, how big a pile, etc.); in a sorites series, the objects are arranged on the basis of the degree to which they possess this property: e.g. the first in the series is the reddest, tallest, biggest pile, etc., and remaining objects in the series are less and less red, less and less tall, less and less big, etc. The arrangement ensure that the following hold:
 - 4a: if $i < j$, $\#i$ is more F than $\#j$.
 - 4b: any object more F than one that is P1 is P1.
5. Adjacent pairs in the series are either indistinguishable, or, if distinguishable, it can seem arbitrary to apply P1 to one member of the pair, yet deny it of the other; likewise for P2.

A sorites series is so called because it is tempting to strengthen (5) to

for all i , P1 and P2 each truly applies to both or neither of $\#i$ and $\#i+1$.

Thus strengthened, the definition leads to paradox, and so to the conclusion that there are no sorites series of objects. By contrast, sorites series are defined in a way that ensures that they are possible; and in fact they are quite plentiful.

To make a *scrambled sorites* one presents each element from a sorites series to a subject in a random order, and elicits a judgment about whether the presented object is P1 or P2. Steps are taken to ensure that the subject has no ready access to her own previous judgments. Once she has made a judgment and it has been recorded by the experimenter, the subject is distracted using familiar techniques, like being asked to engage on a task having no bearing on the topic of the experiment, or being asked to focus on loud music, or whatever. When a new judgment is elicited, the subject should have no clear memory of how she judged any specific previously presented object.

The central question for this paper is: what responses to a scrambled sorites are free from error?

Suppose, for concreteness, that the sorites series consists in thirty color patches ranging from red to orange. In their sorites order, though each shade (other than the first) is less red than the one to its left, there is no detectable difference, or at least no significant difference, between adjacent patches. They are not presented in this order, but in a random order. At each presentation, the experimental subject is required to punch a button marked Y (for Yes) just if she judges the patch to be red, and otherwise to punch a button marked N. Each judgment is made independently of earlier ones, for these have been displaced from memory. What would it take for a collection of such judgments to be free from error? The question requires an apriori answer. It is entirely independent of what real subjects would be likely to do under these imagined experimental conditions.

Let's suppose that once the experimenter has all the responses, she rearranges them, so that they are no longer in the order in which they were made, but rather in the order corresponding to the sorites order of the objects; this constitutes what I shall call the "soritical" series of judgments. In this series of judgments, the first two or three, at the very least, will be Ys, the last two or three, at the very least, will be Ns. I claim that there are just three ways for the soritical series of judgments to involve error:

1. An N corresponds to a red patch.¹
2. A Y corresponds to an orange patch.
3. An N occurs to the left of a Y.

A response of the third kind counts as an error since, given that according to 4a any two patches differ in degree of redness, this response would violate an instance of 4b:

(*) Anything redder than something red is red.

If (3) occurs, then there are n and j such that:

- i. $\#n$ is judged not to be red and
- ii. $\#n+j$ is judged to be red and
- iii. $\#n$ is redder than $n+j$.

Given that (*) is not to be challenged, and that (iii) follows from the definition of a sorites series (in particular, 4a), a type-(3) response ensures that one of the judgments in (i) or (ii) is in error.

Suppose that all three kinds of error are avoided. The soritical series of judgments will consist of a string of Ys followed by a string of Ns. The subject's judgments would in some sense mark a boundary: a last Y immediately followed by a first N, and with no subsequent Ys. There will be many such series, differing just in the position of the first N.² Any such series of judgments will be called a "perfect score".

Apriori reflection on scrambled sorites shows that only by making a perfect score can we avoid error. Making a perfect score involves making a cut: judging a patch to be red that is adjacent to a patch judged to be orange, and so not red. Only by making a cut can error be avoided. The conclusion is that vagueness involves many permissible boundaries or cut-offs. This runs counter to a familiar tradition, according to which vagueness is characterized as absence of cut-offs.

¹ Expressed more carefully: for some i , the i th object in the sorites series of objects is red, and the i th judgment in the soritical series of judgments is N.

² Epistemicists will think there is exactly one such series, though we can never know which it is. This view has to be set aside in this paper, though it concurs with one main point: vagueness does not preclude sharp cuts.

§2. Incoherent scramblings?

In this section, I'll consider objections to the way in which scrambled sorites have been set up.

First objection: there may be no fact about whether errors of types (1) and (2) have been made, since it may be vague, concerning some patch, whether or not it is red. Hence it may be vague whether a score is perfect or not. This is just what we would expect from vagueness, and it blocks the argument to the conclusion that vagueness is consistent with sharp cut-offs.

One way to respond would be to confine errors of types (1) and (2) to clear cases. The errors consist in responding with N to a *clearly* red patch, or with Y to a *clearly* orange patch. But the original unease may remain: "clearly red" is vague, so there will still be series of judgments concerning which it is vague whether they are guilty of type-(1) or type-(2) errors.

A better response is to challenge the claim that vagueness in what constitutes a perfect score undermines the thesis that vagueness is consistent with sharp cut-offs. Given the constraints of the experiment, every soritical series of judgments consists of a string of Ys and Ns. Throw away all strings guilty of type-(3) error, since these are not candidates for perfection. It may be that, concerning the series of judgments that remain, there is vagueness about which count as perfect because there is vagueness about which avoid errors of types (1) and (2). But this does not affect the conclusion. All the candidates for perfection will consist of a string of Ys followed by a string of Ns. We can be sure that all perfect scores have the sharp cut-off feature, an unreversed switch from Ys to Ns, whichever exactly the scores are that meet the further requirement of avoiding errors of types (1) and (2).

Second objection: the set-up for scrambled sorites distorts the nature of vagueness, because subjects are not allowed to answer in any way except "Yes" or "No". To move us closer to reality, we should allow at least for a third option, corresponding to any of "don't know", "no opinion", "would rather not answer", etc.

Adding further options for responses makes no difference to the point about cuts. Suppose the subject is given a third intermediate response option, D (described in whatever way you prefer – e.g. “Don’t know”). Then a perfect score will consist of a row of Ys followed by a row of Ds followed by a row of Ns. The reasoning is structurally as before, with the additional principles

(**) Anything less red than something judged “Don’t know” should not be judged red.

(***) Anything more red than something judged “Don’t know” should not be judged not red.

These principles have as much weight as the indisputable (*). Together with it, they ensure that only by making a perfect score – Ys followed by Ds followed by Ns, with no D to the left of a Y or to the right of an N – can a subject avoid error. Using three responses means that avoiding error involves making two cuts – and similarly if more responses are used. The only constraint is that the responses fall into the kind of order characterized by the three asterisked principles.

Third objection: given the following principle

EFJ: Error free judgments are true

it would be inconsistent to suppose that there are distinct perfect scores (i.e. distinct collections of error-free judgments). For consider two distinct perfect scores, S1, S2, where according to S1, #11 is red and according to S2 #11 is not red. By definition of “perfect score”, both of these judgments are error free. But if both were true, it would be true that #11 is red and that it is not red. This shows that we must either accept that something has gone wrong with the account given of perfect scores, or else accept that EFJ is false.

We already have reason to suspect that EFJ is not in general acceptable. Consider a case in which two gourmets, of high intelligence and fine powers of sensory discrimination, are exposed to liquorice. Their carefully considered judgments are intuitively error free. One judges that liquorice is tasty. The other judges that it is not. If EFJ were true, we would have a contradiction. But intuitively we have merely a case of

“faultless disagreement”.³ Even those who are not happy to accept this diagnosis would not be well advised merely to take EFJ as a premise. At a minimum, it would need independent support in a debate of this kind.

EFJ is certainly initially attractive. If there are cases in which it fails, what is interesting is to explain how this can be so. In the example of liquorice, the failure of EFJ can throw light on the nature of taste. In a similar way, I will argue later (§5) that its failure in the case of vagueness guides us to interesting features of that phenomenon. Before addressing that question, two sections discuss the polar opposite of scrambled sorites: forced march sorites.

§3 Forced marches: semantics versus epistemology

In a forced march, a subject is exposed to the elements of a sorites series in their soritical order, with memory of at least some previous judgments. At a minimum, when making every judgment except the first, both the previous element and the judgment she made concerning it will be present to her mind. For example, she may be shown two color patches. One she has already judged to be red, and she knows she has done so. The other appears indistinguishable (or at least not significantly different). Under these circumstances she naturally feels some compulsion to judge that the other patch is also red. If we press forward with demands for judgment, we can expect that she will:

1. overextend the application of the predicate (relative to judgments she would have made in other contexts): that’s the inertial effect;
2. eventually “flip” – she will judge a patch not to be red even if indistinguishable from one she knows she has just judged to be red;
3. display hysteresis: that is, the probability of a flip occurring at a certain point in the series will be sensitive to the inertial effect in the direction in which she is being force-marched;

³ See Kölbel 2003.

4. be anxious in the borderline area, worrying about whether she is making the “right” call.

In this section, I contrast two explanations of these phenomena. The one I prefer supplies a purely epistemic account, from which it follows that forced march sorites reveal nothing of interest about the nature of vagueness. The other is a kind of scoreboard contextualism, going back to Kamp (1981), which I’ll label “semantic”. According to this explanation, which things are in fact red depends on the context in which the question arises, in particular on what prior judgments about redness have been accepted in that context.

Scoreboard contextualism can be illustrated by an example from Lewis. In a context in which participants in a conversation have accepted “France is hexagonal”, and so have “added this proposition to the scoreboard”, the standards for descriptions of country shapes are set as loose. It follows that, in the context, “Italy is boot-shaped” is true.⁴ In forced march series, the salient features of the context are the subject’s prior judgments. The scoreboard contextualist proposes that, up to a point, accepting a judgment that, for example, some patch #n is red ensures that the judgment “#n+1 is red” is true, if need be by actually expanding the extension that “red” had before. The state of the scoreboard at a time – the judgments accepted by parties to a conversation at that time – can affect the extension of the words being used.⁵

At some point in the forced march, the subject can no longer bear to apply the predicate. At that point, the context becomes “incoherent”: it may “assign opposite truth values to one and the same sentence” (Kamp 1981: 252). The explanation points at a thesis about the essence of vagueness in a predicate: it consists in the context-sensitivity of its extension to the judgments in the scoreboard, a sensitivity that helps

⁴ Lewis uses this example in various writings (1972, 1979). Kamp (1981: 245–6) illustrates the contextual dependence he has in mind by examples of cross-sentence anaphor.

⁵ In Roy Sorensen’s elegant formulation: “Hans Kamp, the founder of contextualism, maintained that the extension of vague words orbits the speaker’s store of conversational commitments”. The founding is best attributed to Kamp 1975.

verify the judgments in borderline cases which the forced-march subject takes herself to be forced to make (until the context becomes incoherent).

On this view, inertia emerges as, more often than not, a *correct* response to the expanding extension of the predicate: inertia is simply dynamic sensitivity to truth. Flipping is explained in terms of the eventual incoherence of the context. Hysteresis can be explained as a combination of inertia and flipping. How the contextualist should explain anxiety is the topic of a separate discussion in §4 below.

Scoreboard contextualism counts as “semantic” in that it says that the extensions of predicates may change with context. A contrasting account of forced march phenomena is “epistemic”, in the sense that it has nothing to say about which things are red, but relies instead on a quite general feature of justification for judgment, which could be summarized in this principle:

(DD) Don't discriminate unless there's a relevant difference.

The accounts agree that, in a forced march situation, remembered past judgments influence subsequent ones. They disagree in that whereas scoreboard contextualism characterizes the influence in terms of a shift in the extension of the predicate, the DD explanation characterizes the influence in terms of a quite general principle about justified judgment.

From the point of view of DD, the reasoning of a subject in a forced march situation might be presented:

I have just judged $\#n$ to be red and I can detect no relevant difference between $\#n$ and $\#n+1$. So by DD I should not discriminate these patches, that is, I should judge $\#n+1$ to be red also.

DD is general, and not even restricted to judgments. For example, it tells you that if two children ask for a biscuit, and you have at least two biscuits, you should give a biscuit to both or neither, unless there is some relevant difference between them. It tells you that if two essays are equally good, you should give both the same grade. It applies in cases where vagueness is not an issue: if figures A and B look alike, and you judge that A is square, then you should judge that B is square.

DD is of course defeasible. If you have only one indivisible biscuit, it may be best to pick a child to give it to, rather than deprive both. There is also room for substantive dispute about what counts as a relevant difference. Does it matter that one child is your own and the other a stranger's? Does it matter that one essay is by a freshman and the other by a senior?

In the forced march situation, DD is defeated by a patch # n that strikes you as clearly orange, even if you can't tell any significant difference between it and a patch # $n-1$ you have just judged to be red. If you follow perception, and judge it to be orange, DD kicks in again at that point, producing an inclination to change your mind about # n , and judge it too to be orange. This instability is a characteristic of vagueness, and is a feature that is to be expected if DD influences behavior as a defeasible but default rule.

Although DD can explain forced march phenomena, it has nothing to tell us about the nature of vagueness. We should not appeal to the standard forced march presentations of soritical series of objects to help understand vagueness. Rather we should appeal to scrambled sorites.

§4. The explanations compared: anxiety; forced marches over precise territory

How should we compare the explanation of forced march phenomena offered by scoreboard contextualism with that offered by DD? Since DD is an entirely general principle, and provides an adequate explanation, scoreboard contextualism's explanation is at best redundant. In this section, I suggest two further points: there is a phenomenon that scoreboard contextualism cannot explain; and forced march phenomena can occur even when vagueness is not at issue, suggesting that forced march phenomena are not well suited to revealing the nature of vagueness.

One feature of forced marches is that the subject is likely to experience anxiety: not existential angst, but an irritated frustration at not knowing how best to proceed, or indeed what the point of the proceedings is.⁶ The scoreboard contextualist will predict

⁶ As Geert Keil has pointed out (personal communication), much depends on exactly how the forced march is presented, and philosophers' accounts are generally not very detailed. One possibility is that the

anxiety when the context is close to incoherent, but has no resources on the basis of which to predict anxiety when the subject is at a much earlier point – for example, no longer in the region of definite reds, but some way from feeling forced to declare a patch orange. Indeed, if the scoreboard contextualist attributes semantic insight to subjects, they should be in a state of high confidence in many of the states in which they experience inertia, for the contextual effect is to ensure that a judgment that in other contexts might have been dubious is in fact true. Yet we can reasonably expect uncertainty and anxiety to set in earlier, and this is what contextualists lack the resources to explain.

DD as such also offers no explanation, but it can very naturally pair with a further principle so that the two together supply just what's needed. We are familiar with "slippery slope" arguments; we know where they can lead; in a forced march situation we can see what's coming; we appreciate that there's no unique right answer. If we have some larger scale purpose that provides a reason to draw a line somewhere, our anxiety will be greatly reduced. In the standard sterile setting of a forced march, we are anxious because, anywhere in the borderline region, a unique judgment is demanded when it's arbitrary which judgment to make. This makes early onset anxiety a perfectly rational state, whereas a scoreboard contextualist must regard it as some kind of aberration: in the early stages at least, context is working on one's side, ensuring the truth of judgments in borderline cases.

subject is presented with the whole series of patches, and then asked to attend to the two leftmost ones, #1 and #2, then to #2 and #3, and so on. In such a case one would expect some sort of anxiety to begin from the start, well before a borderline case has been reached. At another extreme, a subject might be presented with two adjacent panels, and asked to make a color judgment about each in turn. The panels will be arranged so as to make it clear that the panel most recently judged remains visibly the same when the next judgment is solicited: only the color in the other panel is changed. (Judgments relate alternately to left and right panels.) One could insert distracting elements, to help conceal the soritical nature of the arrangement. The crucial feature of forced marches remains: when making a judgment about patch #*i*, both it and the patch #*i*-1 that has previously been judged are present to the subject, and the subject knows this, and knows what she has judged about #*i*-1. In this set-up, anxiety is likely to begin later, since the subject is not directly confronted with the soritical character of the experiment.

Forced march phenomena can arise even in the absence of vagueness in a judgment's subject matter, so that these phenomena are a risky basis on which to build a theory of the nature of vagueness. Guided by a casual glance around the stadium you are in, you are asked to give your best estimate of how many people are there, at a minimum.⁷ Your estimate is correct if there are at least that number of people. You respond "3000". Your interlocutor says: "Given the coarse-grained character of your evidence, isn't the answer 3001 just as likely to be right?" It's hard to suppose that you could rationally think that 3001 would be more likely to be false than 3000, for, if you retain your belief that 3000 is a correct lower limit, the supposition that there are not at least 3001 people entails that there are exactly 3000 people, and you know you are in no position to make a sensible estimate of *exactly* how many people there are. In light of these considerations, you agree that 3001 is as good an estimate. Under these pressures, a series of answers will emerge: 3000, 3001, 3002 ... and so on, until you reach a breaking point, and decline to slide further down the slippery slope. Which numbers n satisfy the predicate "there are at least n people in the stadium" is a precise issue (assuming that we bracket any vagueness relating to persons who are partly in and partly out, or to the outer bounds of the stadium), yet we seem to have replicated many features of the forced march: there is inertia (the tendency to extend the judgment across small increments), flipping (the eventual digging in of toes), hysteresis (asking for a maximum and working downwards would no doubt lead to judgments in conflict with those reached when approaching from below), and anxiety.

Scoreboard contextualism cannot apply in this case, at least not in any straightforward way. Contextualists claim that the fact that the scoreboard contains the judgment that patch # n is red very often makes it the case (by adjusting the extension of "red") that the judgment that patch # $n+1$ is red is true. One analog would be this: that the scoreboard contains the judgment that there are at least 3000 attendees very often makes it the case (by adjusting the extension of what?) that there are at least 3001 attendees. But context simply cannot affect the correct enumeration of the attendees. On the other hand, a pragmatic explanation works fine. By DD, you have no reason to

⁷ The example is suggested by one given by Williamson (1994: 217), though for a different purpose.

prefer 3000 to 3001 as an estimate: no evidence available to you reveals one to be preferable to the other. But DD's guidance is defeasible: for some number n you will see that it is unlikely that there are at least n people in the stadium.

When we think about forced marches over precise territory, it is obvious that the influence of remembered previous judgments is likely to lead to distortions. If our aim is to obtain the best estimates we can of the minimum number of people in a stadium, we would never allow the forced march procedure. This makes it puzzling why one would take forced march sorites as the canonical guide to the nature of vagueness.

I now return to the main theme of the paper: the principle that error free judgments are true (EFJ) fails in the case of vagueness, and understanding why this is so reveals something about the essential nature of vagueness.

§5. EFJ: the mandatory and the permissible

EFJ claims that any error-free judgment is true. This is incorrect for vagueness, because vagueness gives license: it ensures that a judgment can be permissible (given all the facts), and so error-free, without being mandatory, and so without being required by the search for truth. A permissible judgment involves no error, but if it is not mandatory, then there's a sense in which it's not true. To set the view out more fully, let's start with a couple of familiar examples.

Paint shop: the manager of an artist's store has a tier of shelves on which to arrange the tubes of paint he is selling. To display all his tubes, he needs to use several of the shelves. In particular he finds he can get the first 14 of his reds, which he arranges starting with the reddest red, on the top shelf. On the second shelf he starts with the reddest of the remaining paints, and the last tube on this shelf is plainly orange. It is permissible for the manager to label the top shelf "Red" and the second shelf "Orange", despite the fact that the reddest tube on the shelf labeled "Orange" is barely distinguishable in color from the least red tube on the shelf labeled "Red". If a customer complains, the manager can properly respond with a phrase marking the exercise of discretion: "That's what

we call/count as orange here”. If a customer comes in with an empty tube asking for a color match, the manager may correctly respond: “Look among the reds”.

- *Residency*: Sc 3 para 17 of the first draft of the University Regulations stated: “Faculty must live near campus”. This was held to be potentially problematic, and so, after a dozen long and rancorous meetings, the finished version included an asterisk attached to “near”, leading to the following note: “For the purposes of this Regulation, “near” means within 20 miles of the Main Building.” The ruling was justified by the university’s wish to be able to assemble an emergency meeting of all faculty within 45 minutes.

Neither the paint store manager nor the university were required to draw the line where they did; but drawing it where they did was permissible. A local cut-off was established. There was no intention to make a ruling that would hold outside the local context. The manager need have no quarrel with another store that used longer or shorter shelves and so drew the line in a different place. The university need have no quarrel with another institution that offered a different interpretation of “near”.

Context is playing a crucial role in making it sensible to select one rather than another place to make the cut. For the paint shop manager, the relevant feature of context is the length of the shelf. For the university administrators, the relevant feature is the desire to be able to summon emergency meetings quickly. Contexts like these do not function in at all the way envisaged by scoreboard contextualism. These contexts don’t affect semantics; they just make it sensible to select one rather than another point that the semantics have independently determined to be permissible.⁸

Where there is no vagueness, the permissible and the mandatory coincide: it’s permissible to assert that P iff it’s mandatory to assert that P (assuming one is going to address the question whether or not P). Otherwise put: in the absence of vagueness, if one is going to make a judgment concerning whether or not P, there’s just one way to avoid error.

⁸ There are further contextual effects: tall for a child is different from tall for an adult, and *nearby* delivers different distances depending as foot travel or car travel is envisaged. These effects are not what scoreboard contextualism relies upon.

By contrast, when there is vagueness, there are permissible judgments that are not mandatory, as illustrated by the paint shop and the residency requirements. There's more than one way to avoid error. But on the usual conception of truth, not all these ways of avoiding error are ways of attaining truth, since they are not mutually consistent.

This underlies the variability of the cuts in perfect scores on scrambled sorites. Each cut corresponds to a permissible boundary that could be drawn in suitable circumstances. The many perfect scores correspond to the fact that a permissible cut can be made in many different places. The multitude of perfect scores corresponds to the fact that no cut is mandatory.

On this view, vagueness and non-mandatory sharp boundaries are perfectly harmonious features. Indeed, a predicate's vagueness consists in the fact that it has many permissible boundaries. We should therefore reject characterizations of vagueness in terms of absence of a cut-off.

§6 Sharp boundaries, many boundaries, and relativized truth

Theorists often speak of "sharp" boundaries, as if there were boundaries of some other kind. A common claim is that vague predicates are those that draw no *sharp* boundaries. Others, perhaps unclear about what a non-sharp boundary would be, have tried to characterize vagueness in terms of absence of boundaries of any kind (Sainsbury 1990).

The present discussion enables one to make sense of this situation. Let each cut made in a perfect score count as a boundary. Let's say that a predicate has a *sharp* boundary just on condition that it has a unique boundary. Then it is correct to say that there are non-sharp (i.e. non-unique) boundaries, and that this is a characteristic feature of vagueness. We can understand the claim that vague concepts are concepts without boundaries as the correct claim that vague concepts are concepts without unique boundaries. Putting the point positively: the characteristic feature of a vague predicate is that it is associated with many boundaries.

In a similar vein we can understand a claim like "a vague predicate does not have an extension" (for an extension is a set, and so precise) as "a vague predicate

does not have a *unique* extension”, or a claim like “a borderline case is one in which the semantics plus the facts don’t settle a truth value” as “a borderline case is one in which the semantics plus the facts don’t settle a truth value *uniquely*”.

Like reflection on disagreements of taste, reflection on scrambled sorites provides a philosophical motivation for relativizing the notion of truth. When the paint shop owner says a tube is orange, he can be understood as saying something true relative to his shelving arrangements. This is why he has no quarrel with another merchant’s different arrangements: the arrangements are in a sense consistent. When the university tells a faculty member he must relocate, since his home address is 23 miles from Main Building, it’s important to be able to show that he has violated the Regulations: relative to them, it is not true that the faculty member lives near the university. Other standards, while intrinsically perfectly acceptable, are irrelevant.⁹

Some early supervenionists claimed that vague predicates are defective or incomplete in their meaning (e.g. Fine 1975: 265), and this drew fire from those who wished to attest to the importance of vagueness in our thought and talk. The many-boundaries approach does not see vague predicates as incomplete; rather, the guiding metaphor might be that they are generous or flexible, inviting different cuts for different purposes. Some supervenionists have understood sharpenings as ways of making a predicate precise, and this drew fire from those who wondered how one could properly describe the semantics of a vague predicate in terms of predicates not only different (because precise), but different in exactly the critical respect (e.g. Fodor and Lepore 1996). In the many-boundaries picture, this worry is baseless: it is the vague predicate just as it is that invites the many cuts, potentially registered as the supervenionist’s sharpenings. They are not products of some process of revision, but simply reflect permissible uses¹⁰.

⁹ An alternative approach would be to retain an absolute conception of truth, and regard a non-mandatory ruling as effecting a local semantic change (thanks to Geert Keil for reminding me of this). Which approach to adopt is a substantive issue, beyond the scope of this paper.

¹⁰ This paper emerged from a class given jointly with Hans Kamp at NASSLLI 2012. Many thanks to Hans for help and comments at every stage. We plan a joint paper on the various ways in which context relates to vagueness. I would also like to thank the following for valuable comments: audiences at

Afterword: linguistic and metaphysical vagueness

This paper maintains that a vague expression is one for which there is a divergence between its mandatory and its permissible use, captured by its association with many permissible but non-mandatory cut-offs. This property is evidently one that only an expression could have. One open question is whether there is a generic notion of vagueness which is potentially applicable to both linguistic and non-linguistic entities. Identifying such a notion would be one way of vindicating the thesis that there is such a thing as “metaphysical” or “ontic” vagueness.

A *formal constraint* on developing a notion of metaphysical vagueness is that it must issue from a generic notion of vagueness of which linguistic vagueness is one species and metaphysical vagueness another. If this constraint is not met, it will be reasonable to doubt whether the alleged metaphysical vagueness is really *vagueness*; whether “vagueness” is being used without equivocation in the two cases.

A *material constraint* on developing a notion of metaphysical vagueness is that it should exclude the kinds of metaphysical indeterminacy that are intuitively different from vagueness, for example the indeterminacy of some quantum phenomena or of an open future.

To illustrate the impact of these constraints, consider the view that the generic notion of vagueness is fixed by absence of determinate fact (see Merricks 2001: 145). More particularly, generic vagueness amounts to there being an object and a property such that there is no determinate fact concerning whether the object possesses the property. This meets the formal constraint, delivering the specific notions in a straightforward way. Metaphysical vagueness is the case in which the object is not a linguistic one, and the property is not a semantic one. Linguistic vagueness is the case

NASSLLI 2012, at the LOGOS conference *Departing from Sainsbury*, Barcelona 2012, and the Editor of this volume.

in which the object is linguistic and the property semantic (for example the object is the word “red” and the property is, for some object *o*, *applying to o*).

This suggestion, however, fails to meet the material constraint. On some views about the future, there is no determinate fact about how it will be, yet this is not a matter of vagueness. Likewise, some quantum phenomena, to which vagueness is irrelevant, may exhibit this generic absence of determinacy.

Another suggestion at least arguably meets the formal constraint and does better on the material one. This is the idea that generic vagueness is a matter of borderline cases. Applied specifically to predicates, vagueness takes the form of there being borderline cases for their application. Applied specifically to quantifiers, it takes the form of there being non-vague open sentences such that the result of attaching the quantifier to them yields borderline cases for truth. Applied specifically to properties, it takes the form of there being objects that are borderline possessors. Applied specifically to spatial objects, it could take the form of there being points in space that are borderline for whether or not the object occupies them. The generic notion couples in a natural way with a specific subject matter to deliver distinct specific conceptions of vagueness, appropriate to the different cases. The approach does not fall foul of anything so far mentioned in the material constraint: neither quantum phenomena nor the open future make room for talk of borderline cases.

However, it is well known that vagueness cannot be correctly characterized in terms of susceptibility to borderline cases, at least not without some more detailed account of what a borderline case is. As various writers have noted, one could introduce a predicate “child*” with the stipulation that it applies to anyone less than 17 years of age and to no one more than 18 years of age. 17-year-olds are borderline cases. Yet intuitively “child*” is not vague.

Linguistic vagueness, as described in this paper, makes room for many cut-offs, suggesting that it would be a mistake to seek a generic notion of vagueness in terms of lack of cut-offs. Moreover, a suitable generic notion cannot exploit the notions of permissibility or mandatoriness, for these are inapplicable to metaphysical vagueness. The world is as it is independently of conventions pertaining to our use of language. It is

hard to see how these language-specific notions could emerge in a natural way from a generic characterization.

What is hard to see may be the case; but, given the theme of this volume, it seemed worth setting out how the conception of vagueness this paper proposes may present an obstacle to an adequate formulation of a worthwhile thesis of metaphysical vagueness.

References

- Fine, K., 1975. 'Vagueness, truth and logic', *Synthese* 30, 265-300.
- Fodor, J. and E. Lepore. 1996. "What cannot be evaluated cannot be evaluated, and cannot be supervalued either." *The Journal of Philosophy* 93(10): 516–535.
- Kamp, J. A. W., 1975. "Two theories about adjectives", in E. L. Keenan (ed.), *Formal Semantics of Natural Language*, Cambridge University Press, pp. 123-155.
- Kamp, J. A. W. 1981. "The paradox of the heap", in U. Mönnich (ed) *Aspects of Philosophical Logic* Dordrecht, Reidel: 225–277.
- Kölbel, Max. 2003. "Faultless disagreement." *Proceedings of the Aristotelian Society* 104: 53–73.
- Lewis, D. 1972. "General semantics", in D. Davidson and G. Harman (eds). *Semantics of Natural Languages*, Dordrecht.
- Lewis, D. 1979. "Scorekeeping in a language game", in R. Bäuerle, U. Egli and A. v. Stechow (eds) *Semantics From Different Points of View*, Berlin, Springer-Verlag: 172–187.
- Merricks, T. 2001. "Varieties of vagueness." *Philosophy and Phenomenological Research* 62: 145–157.
- Sainsbury, R.M. 1990. "Concepts without boundaries". Inaugural lecture published by the King's College London Department of Philosophy. Reprinted in R. Keefe and P. Smith (eds) *Vagueness: A Reader*, MIT Press: 251–64.
- Sorensen, R. 2012 "Vagueness", in *The Stanford Encyclopedia of Philosophy (Summer 2012 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2012/entries/vagueness/>.
- Williamson, T. 1994. *Vagueness*. London, Routledge.