Pragmatic Vagueness

by

William Nava

1 – Introduction

Consider a series of monochromatic patches stretched out one after the other. The leftmost patch (p_0) is a paradigm case of appears-red (R), such that every English speaker, competent in all the relevant respects (i.e., knows how to use color words, is not colorblind, is not hallucinating, etc.), would judge that it is red. As we move to the right, the patches become gradually less red and more blue. The rightmost patch (p_n) is a paradigm case of appears-blue. The shift happens so gradually, however, that for each patch in the series, its apparent color is indiscernible from that of its adjacent patches.[1]

The fact that each patch is indiscernible from those adjacent to it is the *tolerance principle* (TP): for any p_i, $Rp_i \equiv Rp_{i+1}$. Starting from Rp_0, we can prove (via *n* applications of TP) Rp_n.

This is the sorites paradox. By stipulation, we cannot reject Rp_0. We cannot accept Rp_n, unless we want to accept that a paradigm case of appears-blue can can at the same time and in the same sense appear red. Finally, rejecting TP implies that there is some p_i such that $Rp_i \wedge \neg Rp_{i+1}$. Seeing as R is a sensory predicate, it is absurd to claim two indiscernible patches could differ with respect to R. In other words, it is absurd to reject TP.

The obvious diagnosis is that this is happening because R is a vague predicate: there is no unique boundary between its extension and anti-extension. It might seem plausible to respond to the paradox by pointing out that in an ideal language, predicates have distinct boundaries and the sorites paradox does not arise. But what is meant by "an ideal language"? If it is a language that does not reflect how we in fact reason, then what makes it ideal?

Even if we accept the notion of an ideal language, are there reasons – besides avoiding the paradox – to think that such a language cannot admit of vague predicates? The vague predicates

[1] This version of the sorites is due to Wright (1976). The fact that the adjacent patches are indiscernible w/r/t/ the predicate makes it a particularly tough sorites to deal with. Much of what I'll argue depends on this indiscernibility. Therefore, if indiscernible sorites were not possible, this would invalidate my overall case. They obviously *are* possible, however: the R sorites can be physically set up, were one so inclined. See Dummett (1975) for another indiscernible-adjacents sorites around apparent location as opposed to apparent color.

in natural language are not a defect. Even if we could somehow convince everyone not to use R, but instead use a version of it with a perfectly precise cutoff at some exact wavelength of light – let's call this predicate R* – we would not want to. R* would not be as useful in most contexts as R is. It is a *useful* feature of R that it cannot apply differently to objects that are indiscernible w/r/t/ apparent color.

More needs to be said about ideal languages. Formal logic is *necessarily* an ideal language, at least to some extent. If it were not, there would be no difference between it and natural language. So, if logic can be an ideal version of natural language to some extent, why can it not pretend (so to speak) that all predicates are precise? Why is that level of idealization not allowed?

We now face broader questions about the role of logic and its relationship to language. This is not surprising. How we deal with the sorites paradox is contingent on such questions as: Just what is language? What aspects of it should a formal study of logic represent?

In this paper, I take for granted a *pragmatic* view of language. By this I mean the view that:

P1) The meaning of a word is given only by the relevant[2] facts about how competent language users use it.

P2) Natural language is an emergent tool, contingently shaped by evolutionary imperatives.

P3) Formal logic should represent language so as to maximize the *usefulness* of the representation.[3]

I wish to show that if accept the pragmatic view:

C1) No *fully* satisfactory solution to the sorites paradox is possible.

C2) C1 accords with our intuitions.

C3) We ought to admit of various logics of vagueness – including fuzzy, paracomplete, and paraconsistent logics – for the different contexts under which vagueness arises.

2 – Avoiding cutoffs: a survey

Timothy Williamson (1992) proposes an epistemic solution to the sorites: for any vague predicate, there exists a precise cutoff between its extension and anti-extension that we, by virtue of our

[2] "Relevant" is, of course, vague. I will not argue for the elimination of vagueness from language, so vagueness is not in principle a problem for my definition of the pragmatic view. In the interest of clarity about what I mean by "relevant", here's an example of an *irrelevant* fact about how we use a word: the time of day during which people most often say "tree". This fact tells us nothing about the meaning of the word.

[3] The conjunction of P1, P2, and P3 is *all* I mean by "the pragmatic view". I am not here committed to an anti-correspondence theory of truth, nor to any other views typically associated with pragmatism.

limitations as humans, cannot identify. Logicians have mounted a number of strong objections[4] to Williamson's counterintuitive view. These logicians usually favor some alternative solution. The problem is, most of these alternative solutions are themselves (usually unintentionally) epistemic – they cannot avoid positing the existence of some mysterious and unknowable precise cutoff. Any that explicitly attempt to avoid precise cutoffs fail to do so without also losing coherence.[5]

The existence of precise cutoffs is inconsistent with P1. A precise cutoff would be a fact about the meaning of a word not given by use. To the extent that no proposed solution successfully avoids cutoffs, no proposed solution is compatible with the pragmatic view.

I will now consider some of the more popular proposed solutions and show that they either fail to avoid precise cutoffs or fail to be coherent. For the sake of brevity, my survey of these solutions is necessarily brisk and somewhat reductionist.

Most proposed solutions are *essentially* three-valued. Though the details vary, supervaluationist, paracomplete, paraconsistent, and fuzzy approaches all work by assigning a *third semantic status* to borderline cases. A third semantic status is any semantic status[6] other than "only and completely true" and "only and completely false". All these approaches implicitly or explicitly divide the vague predicate's extension into three regions corresponding to the three semantic statuses. None of them successfully avoids precise cutoffs between these three regions.

Let's begin with the supervaluationist approach, oriented around the notion of *admissible sharpenings* of the vague predicate: all the precise versions of the predicate that would seem to be reasonably appropriate. A predicate applies if it does under *all* admissible sharpenings, and fails to apply if it fails under *all* admissible sharpenings. In all other cases, the predicate neither applies nor doesn't. While this avoids a precise boundary between the predicate's extension and anti-extension, it creates a new one: the boundary between the most reasonable inadmissible sharpening and the least reasonable admissible sharpening. This is a problem – in Williamson's words (1994), "the admissibility of a valuation is itself a vague notion."

This is, of course, the "higher order vagueness" (HOV) problem. We can iterate HOV. Using the same supervaluationist scheme, we can segment sharpenings into the definitely

[4] See Wright (1994), Field (2003), and Priest (2017).
[5] There is at least one solution – Unger (1979) – that avoids cutoffs and is coherent. It does so by denying that everyday objects exist. But, even if we deny that everyday objects exist, they still *seem* to exist and *appear* to have properties – this is implausible to deny. It's easy enough to construct a sorites around apparent objects, so Unger's nihilism fails to resolve the paradox.
[6] We can think of a "semantic status" as all the facts about the applicability of truth values to the proposition.

admissible, the definitely inadmissible, and the neither admissible nor inadmissible. Naturally, the only way to avoid these segments having precise cutoffs is to iterate HOV again – on this second order HOV, sharpenings are segmented into the definitely neither admissible nor inadmissible, the definitely not neither admissible nor inadmissible, and the neither definitely nor not definitely admissible nor inadmissible.

So long as the supervaluationist can continue to iterate HOV, he can, so to speak, postpone precise cutoffs. In cases of continuous sorites progressions, HOV can iterate indefinitely.[7] However, most sorites progressions are discrete, including the R (appears-red) sorites described above. Every iteration of HOV for a discrete sorites will narrow in on borderline regions, creating smaller borderline regions at the next order. This means that the iterability is not indefinite – eventually the process yields borderline regions with only one member. The precise cutoff thereby created is not eliminable. I know of no supervaluationist approach that avoids this problem.[8]

We might question whether this is really a problem. Is a fourth-order precise cutoff really as counterintuitive as a first-order one? This is perhaps a matter of subjective intuition. What isn't subjective, however, is that it is just as incompatible with P1. A fourth-order precise cutoff, intuitive or not, would be a fact about the meaning of the predicate not given by use. Furthermore, no matter how abstract or removed from everyday use a fourth-order cutoff may be, it nonetheless distinguishes between adjacent members. For cases like the R sorites, in which adjacent members are indiscernible with respect to the predicate, that distinction is unavoidably unmotivated.

Paracomplete ("gappy") approaches, while different from supervaluationism in many respects, face the same higher-order cutoff problem. The gappy take (like supervaluationism) tends to work with a "definitely" operator: thus, borderline cases are neither definitely true nor definitely not true.[9] Propositions that apply the predicate to these cases bear semantic value *I* for "indefinite".[10] Naturally, a precise cutoff pops up between true and indefinite and between false and indefinite.

Michael Tye (1994) resists this by making the metalanguage vague.[11]

---

[7] What happens at the limit is a matter of controversy. It is not clear that supervaluationism can avoid precise cutoffs even then (see Williamson 1994). However, I will not cover that here, since discrete sorites suffice to make my case.
[8] Fine (1975) provides a supervaluationist account of HOV. Whatever its merits, it doesn't eliminate higher-order precise cutoffs for discrete sorites.
[9] "Definite" and "determinate" tend to be used interchangeably in the literature. I will stick to "definite".
[10] "The third value here is, strictly speaking, not a truth-value at all but rather a truth-value gap." (Tye 1994)
[11] Hartry Field's gappy solution (2003) differs from Tye's in important respects. However, none are relevant to the present critique. The argument made here against Tye's approach can easily be adapted to work against Field's.

In claiming that it is not true that there are sharp transitions between true and indefinite statements and indefinite and false statements…I am not thereby claiming that the predicate "is true"…is extensionally vague. For if "is true" is extensionally vague then it follows that the set of true sentences has borderline members. This requires that there be sentences which are such that it is neither true nor false that they are true. And this, in turn, requires that there be sentences that are neither true nor false nor indefinite.[12] *I maintain that it is not true that there are such sentences. So I do not accept that "is true" is extensionally vague…. Of course, in taking this view I am not committing myself to the view that* [*the truth-value*] *predicates are precise. Indeed, it is crucial to my account that they* not *be classified as precise. For if they were then every sentence would be either true or false or indefinite, and that would…run counter to my claim that it is indefinite whether no statement of the form "M_n is true" is indefinite. Rather my view…is that they are vaguely vague: there simply is no determinate fact of the matter about whether the properties they express have or could have any borderline instances. So, it is indefinite whether there are any sentences that are neither true nor false nor indefinite.*[13]

This approach certainly avoids precise cutoffs – the question is whether it does so coherently. Consider a sorites progression with 100,000 members. Tye would have to say that there is no fact of the matter regarding which member is the first indefinite member. Indeed, there isn't even a fact of the matter as to whether there are or aren't any members that are neither true, false, nor indefinite! This can't be right. To see why, consider a rather coarse sorites with only five members. Here are two ways the truth-values might turn out:

1: True	1: True
2: True	2:
3: Indefinite	3: Indefinite
4: False	4: False
5: False	5: False

[12] Because both "false" and "indefinite" would qualify as "not true."

[13] Paraphrased slightly for brevity.

…where “   ” means any semantic status that isn’t one of the three values. Of course, Tye denies that there is any “   ”. However, he insists that there can be no fact of the matter regarding whether a proposition bears one of the three values. If this were so, then it *wouldn’t* bear one of the three values – it would have a fourth, “no fact of the matter” status – what I call “   ”. We can iterate, and say there is no fact of the matter as to whether there is a fact of the matter as to whether a proposition bears one of the three truth values. This doesn’t change anything: it’s still not the case that it bears one of the initial three, so it counts as “   ”. The relevant possibilities for semantic status are four: true, false, indefinite, and *anything else*, where “anything else” includes all orders of “no fact of the matter”. It follows that the relevant possibilities for the combination of semantic statuses that the five items can take are finite. Each possibility either involves every member bearing one of the initial three values or doesn’t – there simply is no third possibility, because any such possibility would in fact be a case of *not* every member bearing one of the three. Tye’s approach is therefore incoherent when applied to this five-member sorites. It is not possible for it to be indefinite whether all members bear one of the three proposed truth values.

Though unusual, there is nothing in principle unsoritical about a five-member series. Just as the number of possible truth-value arrangements for this five-member sorites is finite, so it is for a 100,000-member sorites. As a result, Tye’s vague metalanguage approach to avoiding precise cutoffs is incoherent for discrete sorites.

The paraconosistent approach is the mirror of the paracomplete. Where the latter allows for truth value gaps at the borderline region, the former allows for truth value gluts: cases where the predicate both applies and does not apply. “Both true and not true” is paraconsistent logic’s third semantic status. Predictably, a precise cutoff exists between the *just true* and *both* regions and between the *both* and *just false* regions. It isn’t clear how paraconsistency can avoid this.

The fuzzy logic solution proposes degrees of truth corresponding to the continuum between 0 (completely false) and 1 (completely true). Applying this to the R sorites, $<Rp_0> = 1$ and $<Rp_n> = 0$. Borderline cases might have values such as 0.5, 0.4327, and $\sqrt{2}/2$. This, again, creates three truth-value categories: 1, 0, and in-between. Moving down the R progression, there must be some p_i such that $<Rp_i> = 1$ and $<Rp_{i+1}> < 1$. In other words, there must be a first member of the series for which the predicate does not *completely* apply. Wherever there is a distinct first member, there is a sharp cutoff.

Though they don't fall into the "three semantic statuses" approach, it's worth saying a word about contextualist solutions. I will discuss only one: Diana Raffman's "homuncular" analysis (1994). Raffman asks us to imagine the faculty that judges our perceptions of vague predicates as being like an elastic band. As we move down the sorites progression, our "appears red" band stretches out more and more. Finally, at some point, we exhaust its elasticity, and a new band takes over – perhaps the "appears purple" band. Thing is, the bands' ranges of application overlap. If we, now using the "appears purple" band, move backwards down the series, the next few patches – which previously appeared red when we were using that band – now appear purple, because that is the band now activated. Whatever the merits of this analysis, it does not do away with cutoffs. There must still be a precise point at which our "bands" suddenly reach breaking point and must switch. There is no obvious justification for this, especially in the case of progressions where adjacent members are indiscernible.

3 – In defense of cutoffs: a survey

Where some logicians avoid the cutoff, others embrace it. A degree theorist might motivate this sharp cutoff as follows. Take the collection of *all* objects to which some vague predicate can apply. It should be possible to order it according to the predicate's applicability. Even in borderline cases, we can usually tell for which of two members of a sorites the predicate applies more. For example, though 80 grains of sand and 100 grains of sand may both be borderline cases of "is a heap", it's clear that 100 grains of sand form *more* of a heap than 80. Once we've ordered the series from most-applies to least-applies, we call only the first member, and those members *identical to it with respect to the predicate's applicability*, "completely true" (i.e., truth value: 1). The first item in the ordered series that deviates *at all* from that first member in any way relevant to the predicate's applicability no longer instantiates the predicate *fully*. And so, though it represents a cutoff, it is a motivated one.

There is a problem: vague predicates don't usually have a clear "most applicable" case. Is there an exact shade of red which is redder than all other shades? Is there an exact number of grains of sand that make more of a heap than any other number (given that *too many* grains presumably make a mountain and not a heap)? The answer seems to be an obvious "no". This issue aside, we need only remember the R sorites to see that this proposal is untenable. Whatever shade of red is selected as the paradigm for appears-red, the first member to diverge from it enough to not be "completely red" will be indiscernible to the eye from the previous member of the series.

Zach Weber (2010) adopts a paraconsistent approach and attempts to motivate the resulting cutoffs as follows: *every* member of the borderline region, by virtue of being both true and false, is a cutoff. For Weber, this solves the issue because it is *unique* cutoffs that we find counterintuitive. This doesn't work: the first "cutoff" is a *unique* a cutoff between the cutoffs and the non-cutoffs, so unique cutoffs are not eliminated.

Graham Priest (2017) also bites the bullet on precise cutoffs. He writes that in order to identify a cutoff one need merely "[r]un down the sequence until you can no longer give the answer 'yes'. That is where it is." He qualifies this by adding:

> *Since meaning is not subjective, but socially embedded, a more accurate guide to the cut-off is to take multiple speakers and circumstances, and average out the answers. This will provide a more robust determination. Note that this is not to say that what is so is 'what an average person' believes: it merely reflects the fact that* ***words are our words, and mean what we use them to mean****. Can one find out this meaning by empirical considerations? Of course: this is what empirical linguistics is all about.* (emphasis added)

The phrase in bold sounds a lot like a (perhaps slightly weaker) statement of P1. If correct, his argument would constitute a strong case for the compatibility of P1 with the existence of precise cutoffs. The view seems plausible: perhaps precise cutoffs are a product of our *aggregate* language use – an aggregate that none of us knows but is, in principle, knowable.

As put forth, Priest's proposed method would yield only an approximation; his solution to the paradox implies *precise* cutoffs. In order for his motivation of cutoffs to work, it must also imply precise cutoffs. It can easily be adapted to do so. If taking "multiple speakers and circumstances" provides "a more robust determination", then taking *all* speakers and circumstances provides the maximally robust determination. This should yield a precise answer.

There are two problems with this approach. One is that individual speakers' responses to the very same sorites will vary (even if only slightly) along irrelevant contextual parameters. To see why this must be, consider the R sorites. There is no difference between adjacent patches that is relevant to the predicate. Therefore, someone's choice of some patch over an adjacent one can only be based on pure chance or on factors that are *irrelevant* to the predicate, such as time of day, current state of mind, predisposed bias in favor of right over left, etc. In other words, the location of the precise cutoff – and therefore the full meaning of the predicate – will depend on factors that have nothing to do with what we normally take the predicate to mean or what we use it for. This violates the relevance requirement of P1.

Priest's proposal faces another, tougher problem: the domain of all competent language users is vague. Someone who knows only five words of English certainly doesn't count. So where's the cutoff? There is obviously no answer – neither for humans nor for omniscient beings. As a result, the proposed precise aggregate does not exist.[14] This invalidates any justification for precise cutoffs that makes the meaning of the predicate a function of its aggregate use – since that is a P1 requirement, this invalidates all justifications for precise cutoffs w/r/t/ the pragmatic view.[15]

Priest makes another argument in favor of precise cutoffs: the "forced march" thought experiment shows that they are unavoidable. Think back to the R sorites. Have someone answer the question "does this appear red to you?" while looking at each member of the series, one after the next. At first the answer will be an easy "yes". But at *some* point the person must give a different answer. The first different answer that isn't a "yes" – "I don't know", "both yes and no?", "uhh", or whatever – is where the cutoff is.

That cutoff, notably, is a fact of how the person *uses* the predicate – promising for compatibility with P1. However, in its current form, this argument only shows that cutoffs must exist in *individual* sorites runs. To fully convince, it would have to be extended to language users as a whole. The fact that the domain of language users is vague makes this impossible to do precisely.

There is another crucial problem with the argument: which response is the first different response is vague. Is a *slight* hesitation before saying "yes" a *different* response? If not, is "uhhh….yeah?"? Or, if a slight hesitation *is* different, how slight is slight? If we make "different response" *too* sensitive, the second response will almost definitely come out as the first different

---

[14] Priest's approach calls for an averaging of all responses. It's possible to agree with Priest's general approach but not with averaging as the correct method. Why not take the furthest answer from the paradigm as the the cutoff? Or the mode? The latter option is intriguing because it avoids the present criticism of Priest's approach: the fact that "all competent users" is vague shouldn't affect the mode. However, it suffers from a different problem. Members adjacent (from both sides) to the mode are likely to be have *almost* as many hits as the mode. It would be ad-hoc and counterintuitive to place the cutoff between two adjacent and highly popular responses.

[15] This includes Williamson's: "A slight shift in our disposition to call things 'thin' would slightly shift the meaning of 'thin'. On the epistemic view, the boundary of 'thin' is sharp but unstable." (1992)

response. In most cases, certainly including the R sorites, that's too soon. What the right balance is is, of course, vague. This vagueness precludes the *precise* cutoff Priest's system implies.

It's tempting to adapt the forced march scenario to avoid this: let's make the subject *type into a computer* either "yes" or "no" (the computer is insensitive to response time). But all this does is introduce an irrelevant set of parameters that amounts to how the respondent deals with being forced to choose between two indistinguishable choices. As noted above, this is incompatible with the relevance requirement of P1.

4 – Conclusion

I have so far argued both that no solution that admits of precise cutoffs can be consistent with P1, and that no proposed solution successfully avoids cutoffs. What should we make of this?

Consider again Priest's forced march argument for cutoffs. It fails because our attitudes toward the the applicability of the predicate are vague. There really is no fact of the matter about which is the first different response. But now let's think about Tye's paracomplete approach again. My argument for why it fails amounts to the claim that there *must* be a fact of the matter about which is the first different member. So which is it? Must there be a first different member or not?

There's a crucial difference between the two cases. In the case of the forced march, the relevant value is the participant's *attitude*. In the case of Tye's proposed three-valued logic, the relevant value is *semantic status*. While attitude is vague through and through, semantic status *can't* be entirely vague. It is this crucial difference that is the source of the paradox's difficulty.

Consider semantic status first. What any given semantic status is is stipulated by fiat. Indeed, there is nothing more to the meaning of a semantic status than its stipulated definition. It is therefore insensitive to parameters not specified by that definition. In other words, a semantic status is *precisely* what it is defined to be. We can try to define it so as to be vague – as Tye does – but, in doing so, we succeed only in precisely defining something vague. We can extend that vagueness to as many levels of metalanguage as we like, and still, the semantic status is nothing more than precisely what it is defined to be. Because even vague semantic statuses are precisely defined, there is always a clear distinction (i.e., cutoff) between any two semantic statuses.

Quite the opposite is the case regarding attitude. Our attitudes are not stipulated – they are naturally occurring. They are not given by a definition. Indeed, when we try to break down just what is relevant to our attitude toward a proposition, we quickly find that we cannot do so precisely. Furthermore, the number of parameters possibly influencing our attitude is, at least potentially, infinite. So, when we take someone down a forced march, without a doubt *something* about attitude changes after the very first member of the series. Of course, nothing about the attitude that is *relevant to the question* changes in that first move. When does the first relevant change happen? This is incurably vague. Changes happen at every step, and there is no fact of the matter about just which are relevant.

To reiterate: the parameters relevant to attitude in a forced march are vaguely vague and infinite in number. Those relevant to semantic status are precisely as defined (even if defined to be vague) and finite.

There is, of course, one way to make semantic status just as vague as attitude: by using vague language to define it. Tye's method, though it extends vagueness into infinite orders of metalanguage, is a precise method laid out in precise terms. To create a semantic status as vague as our attitudes, we need to describe it and use it as vaguely as we do ordinary language.

But then, isn't this how we *already* talk about truth at almost all times *except* when practicing logic? In other words, isn't the point of logic, in part, to talk more precisely than we normally do?

To see what I mean, consider the following geometric operation: start with any object and transform it into a perfect cube of $.3^n$ the volume of the original. Multiply the new volume by $10{,}001^n$, and move the object that many meters in the direction of $(\pi+n)/5$ radians from true north. Recursively iterate 80,000 times. Given any input for volume and location of the original object, this operation outputs a precise volume and location for the resulting cube.

Now imagine that someone tries to carry this out, starting with an object, say, the size of a dresser. He carefully shaves the object down to a perfect cube of the appropriate volume and methodically transports it the appropriate distance and direction. Let's also imagine this person has the time and resources to carry this out all 80,000 times. What happens? Before long, it becomes impossible to continue with any precision. At first this is for pragmatic reasons, but after enough iterations, physical borders become vague. Once we get small enough, it ceases to make sense to even speak of a perfect cube. Even before we get to this point, the location of the object will diverge greatly from that predicted by the operation. The necessary imperfections in measurement will compound quickly, moving the object far from the expected trajectory.

Is this a failing of mathematics? Ought we adapt mathematics to deal with the vagueness of real-world objects? Of course not. Part of the value of mathematics is that it, so to speak, "pretends" that all objects are precise. When we *use* mathematics, we simply keep in mind the limitations that arise from its precision.

Similarly, logic is precise. This is part of its value. We can incorporate some vagueness into it, but inevitably, we do so precisely. And so, when we iterate logical arguments, we need simply keep in mind the limitations that arise from the precision assumed by the logic.

It's useful here to remember just what it is that logic aims to represent: rules of entailment or valid argumentation. Are these rules precise? If we accept P2, we shouldn't be surprised if, in fact, they aren't. All entailment norms are embedded in natural language, since they have been described – and, indeed, discovered – through natural language. If natural language use has contingently emerged out of evolutionary imperatives, why assume that it must follow precise rules of entailment?[16]

Put another way: mathematics borrows the concept of "object" from everyday experience and then stipulates precise definitions, using precise conceptual constructs like points and lines. This translation procedure causes a mismatch between mathematics and the real world of objects. Similarly: logic borrows the concept of "validity" from everyday argumentation and then stipulates precise definitions, using precise conceptual constructs like semantic value. This translation procedure causes a mismatch between logic (even when interpreted by semantic rules) and natural language argumentation.

Nonetheless, just as precise mathematics is useful despite the fact that the identity of real-world objects is not precise, so precise logic is useful even though ordinary language entailment is not precise.

What does this mean for a logic of vagueness? As far as the pragmatic view is concerned – particularly P3 – the field of logic ought to represent entailment in whatever way is useful, even if the representation is never perfectly precise. The goal against which this instrumentalist approach is measured is the goal of logic: an accurate description of naturally occurring entailment.

[16] This isn't to say that the source of the sorites is that "the world" is vague. Vagueness and precision apply only to representations. "The world" is just what it is. However, entailment is at least as much about rationality (and therefore language) as it is about metaphysical necessity. So, even if there is nothing vague about the (nonlinguistic) world, if the rules of *language* aren't precise, then neither should we expect the rules of entailment to be precise.

The following dialogue may be instructive: I place a borderline case of R in front of someone and ask, “is it red?” He says, “uh, kinda.”

“Well, what do you mean ‘kinda’? Does that mean it’s red to *some extent*?”

“Yeah, it’s sorta red but not really.”

“Well, which is it? Sorta or not really?”

“Well, yes, it’s red to some extent.”

“Okay, well, let’s say ‘extent’ is not allowed. If ‘red’ is a yes/no kind of thing, with no degrees in between, what would you say?”

“Oof, I don’t know. Do I have to pick? Like, it’s kinda red but also not.”

“So both?”

“Okay, yeah, it’s both.”

“Well, what if *that’s* not allowed either?”

“I don’t know – can I say there’s no answer?”

What’s happening in this dialogue is that I keep switching contexts. At first, the context admits of degrees of truth, then it doesn’t; it admits of contradictions, and then it doesn’t. The pragmatic view encourages us to equip ourselves to handle all contexts that could reasonably be expected to arise. In other words, it is *useful* to have more than one logic of vagueness on hand.

Intuitively, most people’s immediate attitude toward vague predicates most closely matches degree theory. However, there are cases when degrees of truth don’t work – such as when a decision has to be made based on whether a predicate applies or not, and the decision cannot be partially made. In such cases, the paraconsistent approach may work, assuming it’s possible to carry out both courses of action under consideration. Or, if it is (also) possible to take neither of the actions, the paracomplete approach may (also) be suitable. In some contexts, we have to deal with a base-level cutoff: no in-between region whatsoever. Whatever the context, since these approaches all need to posit precise cutoffs to be coherent, they never *fully* match our intuitions – as we expect, given P2 and the foregoing analysis. Nonetheless, we do the best we can with what we’ve got.

This is not to say, simply, “let them all in!” Some contexts can conceivably arise and some cannot. Furthermore, there are different approaches to, for example, fuzzy logic. It’s possible that only one of those approaches best matches what’s going on in those contexts where fuzzy logic is the right approach. Even for a pragmatist, there is no reason to keep around a logic that is altogether redundant or counterintuitive.[17] There is still plenty to debate about.

Put another way: to be a pragmatist is to value simplicity less than it is traditionally valued. While simplicity is certainly a virtue for the pragmatist, it is not valuable enough to warrant restricting ourselves to only one logic of vagueness. The power that comes with the ability to treat entailment differently in different contexts is too useful to pass up. What's more, it is useful because it is accurate: alternative contexts of use for scenarios involving vague predicates do in fact arise in ordinary language argumentation.

[17] For example, there aren't any obvious contexts that uniquely call for supervaluationism.

REFERENCES

- Dummett, M. “Wang’s paradox” (1975)
- Field, H. “No fact of the matter” (2003)
- Fine, K. “Vagueness, truth and logic” (1975)
- Priest, G. “Dialetheic vagueness” (2003)
- Raffman, D. “Vagueness without paradox” (1994)
- Tye, M. “Sorites paradoxes and the semantics of vagueness” (1994)
- Unger, P. “There are no ordinary things” (1979)
- Weber, Z. “A paraconsistent model of vagueness” (2010)
- Williamson, T. “Vagueness and ignorance” (1992)
- Williamson, T. *Vagueness* (1994)
- Wright, C. “The epistemic conception of vagueness” (1994)
- Wright, C. “Language-mastery and the sorites paradox” (1976)