

ARTÍCULOS

Vagueness: a Statistical Epistemicist Approach

Jiri Benovsky

RESUMEN

Hay tres concepciones tradicionales de la vaguedad: la primera lo considera como un fenómeno metafísico genuino; la segunda, como una cuestión de ignorancia; y la tercera como un fenómeno lingüístico o conceptual. En este artículo presentaré estas tres concepciones, especialmente la epistemicista y la superevaluacionista, y brevemente apuntaré los bien conocidos problemas a los que se enfrentan. A continuación, examinaré una concepción “estadístico-epistemicista” de la vaguedad que está diseñada para evitar estos problemas. Se trata de una concepción según la cual el fenómeno de la vaguedad se origina en nuestras prácticas lingüísticas; a la vez se insiste en que el significado sobreviene al uso, y que nuestro uso de los términos vagos *sí* produce significados precisos y bien definidos, que ignoramos, con lo que se puede seguir manteniendo la bivalencia.

PALABRAS CLAVE: *vaguedad, epistemicismo, superevaluacionismo.*

ABSTRACT

There are three main traditional accounts of vagueness : the first treats it as a genuinely metaphysical phenomenon, the second as a phenomenon of ignorance, and the third as a linguistic or conceptual phenomenon. In this paper I will briefly present these views, especially the epistemicist and supervaluationist strategies, and shortly point to some well-known problems that burden them. I will then examine a ‘statistical epistemicist’ account of vagueness that is designed to avoid precisely these problems – it is a view that provides an account of the phenomenon of vagueness as coming from our linguistic practices, while insisting that meaning supervenes on use, and that our use of vague terms *does* yield sharp and precise meanings, which we ignore, thus allowing bivalence to hold.

KEYWORDS : *Vagueness, Epistemicism, Supervaluationism.*

I

The phenomenon of vagueness appears usually (but not exclusively) in one of the three following cases: vague objects, vague predicates, and the *sorites*. The first case is the one where an object has imprecise (vague) boundaries – Mount Everest, for instance, is both spatially and temporally vague since it is hard to say exactly where and when it begins and ends (what is its spatial/temporal extension). Everest’s spatial and temporal boundaries are simply imprecise. It seems, at least *prima facie*, that all the ordinary objects we usually quantify over (including ourselves) are vague in this way. The second case arises when it is indeterminate whether an object has a given property or not – thus, it is indeterminate whether a given predicate applies to it or not. If, for instance, Tom has 112 hairs, it is indeterminate whether “bald” applies to him, simply because “bald” is vague, and so it is vague whether a person with 112 hairs is bald or not. Many, but not all, of course, predicates are vague in this way. The third case, the *sorites* argument is easily generated by the use of some vague predicate (like “bald”) or by appealing to any vague object (like “heap” or “mountain”). The argument proceeds as follows:

1 grain of wheat does not make a heap.

If 1 grain of wheat does not make a heap, then 2 grains of wheat do not.

If 2 grains of wheat do not make a heap, then 3 grains do not.

...

If 99.999 grains of wheat do not make a heap, then 100.000 do not.

100.000 grains of wheat do not make a heap.

The argument is valid, using only *modus ponens*. The premises all seem to be true, but the conclusion seems clearly false. This is why the *sorites* is often said to be a paradox.

Any theory of vagueness must deal with the three cases just mentioned. There are three main such theories on the market that I am interested in here: the metaphysical approach, the epistemic approach, and the supervaluationist approach. Let us briefly see what treatment of our three cases of vagueness these views provide.

II

The metaphysical approach simply takes the phenomenon of vagueness involved in all three cases at face value. Everest is a metaphysically vague

object that has vague boundaries. The reason why it is hard to say where/when it begins is that it does not have a precise spatial/temporal beginning. In the same way, it is indeterminate whether Tom is bald, simply because there is no determinate fact of the matter whether Tom exemplifies baldness. Consequently, this approach rejects bivalence since propositions like “Tom is bald” will turn out to be neither true nor false. And this is also how the view deals with the *sorites*: some premises of the argument will simply not be true (while not being false) and the argument will then simply not go through.

Besides the unwelcome fact that this approach in its standard form commits us to the rejection of classical logic, since it rejects bivalence and truth-functionality, perhaps the main charge that is often raised against this view is that it is simply incredible. The world isn’t vague, objectors insist, and while of course such a worry is not enough in itself to refute the metaphysical approach, it is suggested that other treatments of the phenomenon of vagueness should be given. (I shall not discuss here in detail the metaphysical approach; see Barnes (forthcoming) and Williams (2008) for a different defence of this view that takes metaphysical indeterminacy as a primitive phenomenon.)

III

Thus, if one does not find the metaphysical approach appealing, one can rather claim that Everest *does* have a precise spatial and temporal boundary, but that the reason why it is hard to say where/when it begins is that we simply don’t know where the boundary is. So the world isn’t vague, but we are ignorant of the sharp edges that delimit the objects that exist. Similarly, there *is* a fact of the matter as to whether a person with 112 hairs is bald or not, but we simply don’t know this fact. There is a sharp threshold where a person with n hairs is bald and a person with $n+1$ hairs is not. And there is a sharp threshold that determines the exact boundaries of Everest. We just don’t know where these thresholds lie. The *sorites* is thus easily dismissed by claiming that one of the premises is false since there will be a point where one grain of wheat will make a difference between a non-heap and a heap.

As strange as such a view may seem, it has a weighty advantage: it preserves classical logic, bivalence (as well as the law of excluded middle), and truth-functionality entirely. Indeed, according to this view, there are no truth-value gaps. This is easily seen in the case of the *sorites*: all of the premises have a determinate truth-value (true or false), and the conclusion is false because, since one of the premises is false, the argument is simply not sound (while being valid).

This is all very nice, the objector says, but the view is no more acceptable than the metaphysical approach. There are three worries that an epistemicist view must address in order to gain plausibility:

- (i) The view appeals to an epistemic phenomenon of ignorance which is largely mysterious and under-explained.
- (ii) It seems *prima facie* strange (to say the least !) to claim that Everest really has a precise boundary in the very strong sense that if you took just one molecule out of it, it would thereby cease to exist, or to claim that one grain of wheat makes a difference between a heap and a non-heap.
- (iii) There is a worry about arbitrariness that arises from the idea that there is a sharp boundary, for one may always ask “Why *this* particular boundary and not a different one?”

IV

Before turning to what I believe is the right solution for the epistemicist, let us now turn to what is a more popular position: supervenience. According to this view, vagueness is a linguistic (or conceptual) phenomenon. It is not the world that is vague, but our words (or concepts). “Everest is vague” is true because nobody ever gave a precise enough meaning to “Everest” that would determine Everest’s boundaries exactly. It is vague whether some person is bald, because nobody ever gave a precise enough meaning to “bald” that would determine exactly how many hairs one has to have to be non-bald. Given this lack of definiteness of our words, the propositions that involve borderline cases of bald people or mountains are said, as in the metaphysical approach, to lack a truth-value. But many propositions in the neighbourhood *have* a truth-value. The core idea here is to treat the phenomenon of vagueness as *semantic indecision* (see Lewis (1986)). There are in the world many non-vague candidates for being the referent of “Everest” that have sharp spatio-temporal boundaries, and it is just undecided to which one of these candidates “Everest” refers – but were we to take a decision to *precisify* exactly what the spatio-temporal boundaries of Everest are, only one candidate would then be the referent of “Everest” (the candidates for being the referent of “Everest” are often called *precisifications*). What we have then is a situation where our propositions are true/false only under a precisification – “Tom is bald”, for instance, is determinately true/false under any admissible precisification of “bald” (if “bald” is to mean to have at maximum 100 hairs, then Tom is non-vaguely non-bald since he has 112; but, of course,

under another precisification he turns out to be bald). But, as we have seen, the non-precisified proposition “Tom is bald” is truth-valueless. Nevertheless, according to the supervaluationist’s semantics, while a proposition is not true/false, it is said to be *super-true/false* if it is true/false under all admissible precisifications – for instance, it is super-true that Tom is bald if Tom has 0 hairs, since under all admissible precisifications of “bald” he turns out to be bald, and it is super-true that Tom is bald or Tom is not bald since the disjunction turns out to be true under all precisifications of “bald” regardless of the amount of hairs on Tom’s head. At the ‘lower’ level (the level of precisified truths, not super-truths), the *sorites* is then blocked in a similar way to the epistemic approach: under any admissible precisification of “heap”, one of the premises of the argument will turn out to be false, and so the argument will not be sound; and at the higher ‘super-level’ (the level of super-truth, which is the one that matters to the supervaluationist) it is blocked in a way that rejects bivalence (and is, in this respect, similar to the metaphysical approach): since, at this level, many of the premises of the argument will turn out to lack a definite (super-)truth-value, the argument will not go through either. This does *not* mean, defenders of supervaluationism will say, that there is a sharp threshold for being a heap, since supervaluationism doesn’t say that any one particular premise of the *sorites* is false, but merely that some premise is (super-)false.

But, the critics say, let’s be honest – this will not do. I tend to agree: it doesn’t seem so clear that the supervaluationist can avoid being committed to the existence of a sharp threshold between a heap and a non-heap. As we have just seen, at the ‘lower’ level of truth, there *is* such a threshold under any precisification, and so it is hard to see how, at this level, supervaluationism is any different from epistemicism in this respect. And, we also encounter a commitment to a sharp threshold at the super-level, which is easily seen in Tom’s case, since under supervaluationism it turns out to be super-true (because true under all precisifications) that there is an n such that a person with n hairs is bald and a person with $n+1$ hairs is non-bald. And what does it mean to say, as the supervaluationist does, that there *is* a threshold but there is no *particular* threshold? This amounts to claim that there can be true existential statements that have no true instances, which seems counter-intuitive (to say the least!). Is it not then more plausible to claim that there is a threshold, and that there is a particular threshold that we are ignorant of? However counter-intuitive this epistemicist claim might be, it is no less counter-intuitive, I would dare to suggest, than the supervaluationist’s claim (but: only at the two conditions specified just below), while it has the weighty advantage to fully preserving classical logic – whereas supervaluationism not only has problems with instances of true existential statements, but it also rejects bivalence.¹

My point here is *not* to claim that the supervaluationist strategy fails. Indeed, there are many alternative logics on the market, some very peculiar but some that are actually very close to classical logic, and the supervaluationist can certainly choose to defend one of these to back her view up. Instead of ‘rejecting’ supervaluationism, what I wish to point out here is that in the trade-off game of weighting pros and cons, epistemicism seems to compete at least as well as supervaluationism, but only if it could gain intuitive acceptability by answering the three charges against it, that we have seen above in §3. Especially, it is in need to explain the nature of our ignorance of the sharp thresholds there are in the world, and it has to defend the claim that there are such thresholds in the first place.

The view I shall now advocate for is an epistemicist view designed to avoid the flaws of supervaluationism by preserving bivalence and thereby accepting the existence of sharp thresholds between heaps and non-heaps, but also to be able to address the three above-mentioned worries that are raised by the ‘standard’ epistemicist view.² This approach will use a statistical rule instead of the rule of supervaluations. In short, the view I propose claims that, first, there *is* a sharp threshold from being a heap to not being a heap which corresponds to a ‘statistical heap’, a notion that will be explained below. Second, the phenomenon of vagueness is a phenomenon of ignorance, but for a very simple and non-mysterious reason: we never usually ‘compute’ exactly the statistics involved in determining the exact nature of the ‘statistical objects’. This has the consequence that while a statement that involves vagueness (“This heap is F”) has a definite truth-value, we can only know it approximately – so it is only approximately true, as far as we can tell. My proposal will be one of a *general strategy*, where I will leave some questions open, and I will not take a stance and choose between several possible variants of this general strategy – indeed, I will argue that some amount of flexibility of the view I propose is actually welcome. I shall first examine and try to make plausible a version of this general strategy that will however ultimately fail (§5-6), but only to give rise to a second version that will succeed by acknowledging the fact that vagueness is a linguistic phenomenon (VII).

V

Take a practical example. If two neighbour peasants are to clearly delimit the borders between two fields that have been in their families for generations so that no record exists to say where the borders exactly are, and so nobody knows it, what they will be likely to do is just to pick the middle (average) candidate of all acceptable candidates for being the frontier between their fields. This also corresponds to our common practices when we are faced with the question of exactly where and when Everest begins, and when we are

at least implicitly aware of the spatio-temporal vagueness involved – we will often simply say: well, approximately there and then. That is, we will simply pick one of the candidates for being Everest that will probably be roughly in the middle of all of the admissible candidates. This candidate has nothing ontologically special or privileged compared to the others of course, but it is intuitively the one we will think about when trying to say with some precision where and when Everest begins. We'll simply pick something like 'the average Everest'. And to make this intuition more precise: we will pick (approximating it as precisely as we are able to) something that we may call '*a statistical object*' that is Everest (let's call it "Everest_s"), which is simply the 'statistical object' that supervenes on all of the admissible Everest-candidates there are ("supervenes on the candidates" means here that it is 'computed' as being the average candidate). Thus, Everest_s has a sharp spatio-temporal boundary, a precise spatio-temporal location, that is the average location of all of the Everest-candidates. As already mentioned, it is no addition to ontology to postulate the existence of a statistical object like Everest_s since it is really nothing more than one of the candidates that are already there; and by claiming that *this* candidate is Everest_s, we are not making it in any way ontologically special, we are just saying that it is the average one.

So here we have a view that provides a bivalent and epistemicist account of vagueness. First, there *is* a sharp threshold from being Everest to not being Everest, and it is Everest_s. Second, we don't know which one of the Everest-candidates is Everest_s, for a very simple and non-mysterious reason: we never usually 'compute' exactly the statistics involved in determining the exact nature of Everest_s. But, in principle, we could, while it may be a practically unachievable task (for human beings). Since an omniscient God would know the exact nature of Everest_s (having made the 'computations'), there would be no vagueness in the world for her.³ But there is for us, and there always will be, since all we can do is to make more and more precise approximations to get to Everest_s, without probably ever being able to achieve it with God-like precision. And this means that while a statement that involves vagueness ("Everest is F") has a definite truth-value depending on whether Everest_s is F or not, we can only know it approximately – so it is only approximately true,⁴ as far as we can tell, and when we refer to Everest, we approximately refer to Everest_s. And this accommodates perfectly well the phenomenology of our beliefs about the truth-value of statements that involve vagueness. (Most of our statements are then, from our epistemic point of view (as far as we can usually know), approximately true/false, rather than true/false. But of course this does not prevent all statements to have a definite non-approximate truth-value.)

So the strategy of 'statistical objects' is a way to have unknown sharp thresholds in the world that are well-explained and non-mysterious, and so the view answers the two charges against epistemicism we have seen above,

while of course retaining its most weighty advantage over supervenience: that of preserving bivalence, and all of classical logic, and that of providing a simple and straightforward treatment of the *sorites*.

Did we dissipate all worries of those who give epistemicism an incredulous stare? Probably not. But since on the present proposal Everest_s supervenes on all of the Everest-candidates, it in a sense *stands for* all of them, and so while it is a non-vague referent of “Everest”, it still does intuitively justice for the phenomenon of vagueness because it takes into account all of the candidates (the borderline zone, or the ‘fuzzy zone’ of Everest’s boundaries): thus, if bulldozers come and modify the nature of some of the Everest-candidates, they will thereby also make changes in Everest_s. I believe that this makes the epistemicist claim for there being sharp thresholds in the world less incredible, and that it improves its already good position in the competition for being the best account of vagueness. But nevertheless, I still understand those who still stare incredulously at it. But one can stare back at the ones who stare since they have to concede that the statement “Everest has precise and determinate boundaries” turns out to be (super-)true under supervenience.⁵ (I will say more about the counter-intuitiveness of the statistical rule in §7 below).

VI

A serious objection seems to arise here: where do the candidates for being Everest_s come from? More precisely, what determines which objects are the candidates for being Everest_s? It seems that there is a problem of circularity. Everest_s’s spatio-temporal boundaries are determined (“statistically calculated”) by the candidates for being Everest_s. But how do we identify these candidates? Presumably, by being candidates for us, that is by being epistemic possibilities for being Everest_s. But how can something be an epistemic possibility for being something if the latter something is unknown to us? How can something be an epistemic candidate for being Everest_s if we don’t know where Everest_s is? Simply put, it seems here that the candidates are determined by the nature of Everest_s (because our knowledge of what the candidates are depends on our knowledge of Everest_s – we simply need to know where Everest_s is in order to be able to pick out candidates for it); while at the same time the nature of Everest_s is determined by the nature of the candidates (and likewise for our knowledge of Everest_s boundaries).

To understand more fully what the objection says, let’s consider the following (flawed) line of response. One could say that the objection makes a mystery of something that is not mysterious at all. The objector raises the question “how can something be an epistemic possibility for being something if the latter something is unknown to us?” But, the answer could go, this is

actually a common situation. In Agatha Christie's novel *Ten Little Niggers*, there are 10 candidates for being the murderer, and of course we do not know which one is the real murderer, but this does not prevent us from having a clear notion of what a candidate is, and which ones are the candidates. And so the same can be said about the Everest-candidates: we do not know where Everest_s is, but this does not prevent us from having epistemic candidates for Everest_s.

However, this answer to the objection would of course not work because the two compared situations are not analogous. The crucial difference is that the murderer's identity is metaphysically determined independently of the candidates, while Everest_s's identity supervenes on the candidates' identities.

But while this response, as such, does not answer the objection, it can easily motivate a (better, but still flawed) answer. We know one thing, and we ignore another. The thing we know is roughly where Everest_s is (we know for instance that it is in the Himalayas and not in the Alps); the thing we ignore is where exactly Everest_s is. So, problems with higher-order vagueness set aside, we know where the candidates are, and this is because they are metaphysically given. We could simply reject here the objector's claim that the answer to the question "how do we identify the candidates?" is that they are epistemic candidates. We could instead claim that they are metaphysically given. So what is being proposed here is that at the metaphysical level, things are given (and precise, non-vague): there is Everest_s, there are the candidates for being Everest_s, and there is a strong relation of ontological dependence between them: they all just come into existence together. Everest_s's nature depends on the candidates' nature, and vice versa. And now at the epistemic level: we know many of the candidates (not all presumably, and with only limited precision anyway, and with the possibility of making mistakes and taking false-candidates for being candidates), this explains why we know roughly and approximately where Everest_s is. But because we don't know exactly the candidates, and also because even if we knew all of them we would not be able to compute the "statistical middle" of them (or we simply wouldn't care to do the computations), we are ignorant of where Everest_s is.⁷

VII

So far, we have seen how the epistemicist strategy could be improved with a statistical tool in order to explain neatly and non-mysteriously of what exactly consists our ignorance of sharp boundaries in the world. But the price paid in §6 above is too high: the claim that the candidates for being Everest are metaphysically given and determine together Everest_s perhaps takes away a mystery about our knowledge (our ignorance) of Everest_s but only to replace it with a mystery about our knowledge of the candidates. For how is it that some-

thing is a candidate for being Everest, if this is a metaphysical matter, independent of us? Let us take an extreme example to make the claim as clear as possible – why, for instance, is the table that is before me not a candidate for being Everest? Or why isn't some closer object to Everest, like the K2 mountain in the Himalayas, not a candidate for being Everest? The point here is that if it is a brute metaphysical matter that some things are candidates for being Everest, we then lack a link between the candidates and our knowledge of them.

What I think that in the end all these considerations hint at is that it simply has to be recognized that the phenomenon of vagueness *comes from us*. The world isn't vague, and there are no brute metaphysical sharp boundaries that we are ignorant of *in the world*. Rather, it is our words and concepts that are vague and imprecise. Vagueness *is* a linguistic (or conceptual) phenomenon. But embracing a linguisticist attitude towards the phenomenon of vagueness does not mean that one has to associate with an unpalatable ally, namely supervaluationism, with which I have noted dissatisfaction at the beginning of this paper (after all, Williamson's own epistemicist position is also a linguisticist one).

So what I would like to propose now is (i) to keep using the statistical strategy while accommodating the fact that (ii) vagueness comes from us, (iii) bivalence holds, (iv) supervaluationism is unnecessary, and so we can wholly retain classical logic and avoid the flaws of supervaluationism.

The rough intuition that the statistical entities strategy is based on, is that whenever we have doubts about where to draw a boundary in vague cases, we have a disposition to draw them somewhere in the middle of the zone of vagueness. Remember the case of the two peasants: I think that the problem of vagueness can be solved very much like this practical one.

The starting point here is one that is shared by many, epistemicists and linguisticists alike: the meaning of our words is determined by the use we make of them. (As the slogan goes: meaning supervenes on use.) The linguistic approach insists upon the alleged fact that our use of vague terms, like "Everest", is *not* such that it gives these terms a precise meaning – in this particular case, it does not define Everest as having precise spatio-temporal boundaries. *This*, however, is not true once we embrace the statistical view sketched above. The story is very simple: just take all⁷ uses of the term "Everest" and use them to 'compute' an average meaning (that will correspond to Everest_s). Even if the actual uses of the given term are vague, averaging all of them will yield a sharp and precise result. This sharp and precise result *is* then the meaning of the term "Everest", a meaning that supervenes very clearly on the actual uses of the term – and a meaning that is, in accordance with epistemicism, unknown to us, since we are simply not able to collect all of the uses of the term and make the appropriate average. And of course, the end of this simple story is that bivalence can be maintained, and the *sorites* easily solved, in the same way it is solved by epistemicism, but without any

mysteries concerning our ignorance of the sharp boundaries of heaps or mountains. The sharp boundaries there are, are not metaphysical facts somehow mysteriously inaccessible to us; rather, it is our words that yield sharp extensions for vague terms.

What about grains of wheat? That is, how does this strategy solve the *sorites* paradox? A *prima facie* worry can be formulated as follows: if you take all the heaps, or all acceptable candidates for being heaps, say for instance all mereological sums of grains of wheat between 9 grains and 50 billion and 9 grains, the average of those will be 25 billion and 9 grains, which should then be the sharp boundary for heaps – the point where a heap turns into a non-heap – which is clearly wrong (since, for instance, 1 billion grains clearly make a heap). But this would be to forget about all the negative and indirect uses of “heap”, like the negative and indirect uses of “Everest” mentioned above in footnote 7. Think again of the case of the two peasants. One can say: “This is Tom’s field” and “This is John’s field”. By saying, “This is John’s field” one relevantly says “This is not Tom’s field”, but since the boundary between the two fields is vague, there is a vague zone of overlap between the two fields, and this is where I say that a reasonable and natural thing to do is to draw the border in the middle of this zone. Similarly for Everest: it is by taking into account all direct, indirect, positive, and negative uses of “Everest” (“This is Everest” and “This is not Everest”) that it is possible to get to ‘compute’ the precise and sharp meaning of “Everest”. And the same goes for heaps, cases of baldness, or redness, or any other vague terms: the precise meaning of the word “heap” does *not* supervene just on acceptable candidates for being a heap, but on acceptable uses of the word “heap” in all sorts of statements (or thoughts, see footnote 7), including statements like “This is not a heap” when pointing to a car, when pointing to 2 grains of wheat, or when pointing to 100.000 grains of wheat arranged in an unsuitable way (say, too scattered, for instance). In short, the ‘computation’ is not just the average amount of grains in a heap, rather it is a linguistic rule that takes all uses (even indirect, implicit, negative, or merely thought) of “heap” and yields a precise meaning of the term, that is, the average one. This is, of course, much easier to do and to conceive in any *particular* case of a particular heap, like Everest (that is, after all, nothing more than a big heap of rocky bits); but it works in the general case in a similar way.

Perhaps the strongest objection to this view is that it simply seems to be too arbitrary to be credible (see Heller (manuscript)). Why take an average rather than a weighted average? Why not situate the meaning of the term at an ‘average’ of 67% rather than 50%? What makes 50% special?

I think that there are two kinds of attitudes available to face this worry. First, one could perhaps make good sense of accepting that something more complicated than a straight average of all uses of the term can be used. As a start, only *competent speakers*’ uses should perhaps be allowed to count. Per-

haps my use of “Everest” in some of my dreams should not count; perhaps at least some of the uses of “Everest” in philosophical debates should not count; perhaps uses of “Everest” by very small children should not count. In any empirical research, when doing a statistical survey or when doing any testing of a certain group of subjects, some of the subjects (or some of their answers) are simply disregarded,⁸ perhaps because their answers are too erratic, or perhaps because they are too tired, not sufficiently concentrated, and so on. But not only does it seem very reasonable to help ourselves with a notion of a competent speaker, and so disregard some uses of “Everest”, but it also seems reasonable to think that some competent uses are *more* competent than others. For instance, inhabitants of Himalaya are perhaps more competent with respect to Everest’s spatial boundaries than inhabitants of London – and so, perhaps the average that determines the meaning of “Everest” *should* be a weighted average. Of course, deciding how to weight the average would probably be a very difficult task – and a task that would again yield the arbitrariness worry.

Another way to put these worries is to ask what kind of uses count to determine the meaning of “Everest”, that is, to ask on which uses the ‘computation’ is to be based. There are various possibilities: (i) either only *token* actual uses count – that is, we only take into account uses of “Everest” that people actually utter (in different languages, of course) – or (ii) linguistic *dispositions* to use “Everest” also count, or (iii) we *do* only count actual uses (so, not dispositions) but rather than considering all tokens of utterance of “Everest” we concentrate on *types* of uses. The latter strategy’s motivation is grounded in the idea that if someone were to sit at the bottom of Everest and were to repeat continuously “Everest, Everest, Everest, ...” by pointing towards Everest, then perhaps these token utterances should not be taken to have too much effect on the statistical calculation. And as we have seen above, there are also some special kinds of uses of “Everest” concerning which it is unclear whether they should count or not. If a three-year-old child says “Everest” when pointing at a kitchen table, if someone who suffers from a mountain phobia disorder says “Everest” when pointing towards a Rorschach-test-like picture, or when a philosopher keeps saying “Everest” while giving a talk – then we should probably disregard these tokens and claim that they shouldn’t count in the overall statistical calculation that determines the meaning of “Everest”.

The two general questions that arise from all the considerations above are: if meaning supervenes on use – which use is it, and how does it supervene? Both questions can be put as questions about weighting: if some uses are to count ‘more’ than others, it is then a ‘weighted average’ rather than a ‘straight average’ that should be done, while determining how the weighting should be done would also tell us how the overall statistical meaning of “Everest” supervenes on uses of this term. It seems to me that it is in the nature of any natural language that it will always remain unclear, for any term, which uses of it count as proper and ‘better’ than others (in the sense men-

tioned above). The lack of a good and precise answer to the rightful questions that arise here thus comes not from the account I wish to submit, but rather from any account that wants to claim that meaning supervenes on use. Even if this is right, it of course, frustratingly, does not answer the questions, but it leaves open various possibilities (depending on the various stances one can take with respect to (i), (ii), and (iii) and with respect to the subsequent questions) which to my mind is a theoretically virtuous move, since it shows that the view I am advocating does not constrain one to defend a strong and probably very controversial and difficult-to-defend view about which uses should count as meaning determination and how. Instead of seeing this lack of a precise account as a defect, I would like to suggest that it actually shows the flexibility of the general ‘statistical’ strategy that could help the epistemicist and that is compatible with various answers an epistemicist might provide to the question about how meaning supervenes on use. The fact that there is no good and precise answer here is rooted in the very nature of natural languages, and since the phenomenon of vagueness itself is a linguistic phenomenon, not only is such an ‘incomplete’ account not to be seen as a defect, but rather it is to be understood as a reflection of the pervasiveness of the phenomenon of vagueness in all natural languages.

To conclude, I think that we should adopt a humble attitude towards arbitrariness: while acknowledging that the ‘statistical move’ *is* arbitrary, it should be remarked that it is *not unacceptably* arbitrary. In the case of the two peasants, for instance, we simply have the tendency to take a straight average whenever we need to precisify vague issues, and even if we recognize it as being to some extent arbitrary, such seem to be our common practices. The view I tried to defend in this paper, while being genuinely epistemicist, agrees with the (good) intuition behind the linguisticist view that vagueness is a phenomenon that comes from us, from our language, and from a lack of precise decisions about what our words mean. But instead of appealing to a rule of supervenience that causes more trouble than it is worth, the view I have examined here claims that there is a precise and sharp meaning that supervenes on use, and that allows bivalence to hold. I think that, once the advantages and drawbacks of the statistical approach, of supervenience, and of ‘standard’ epistemicism are put in the balance and carefully weighted, the price of the statistical approach (that is, a certain (acceptable) amount of arbitrariness and imprecision) is worth paying – or at least a serious buyer could/should consider paying it.⁹

*Department of Philosophy
University of Fribourg
Av. de l'Europe 20, 1700 Fribourg, Switzerland
E-mail: jiri@benovsky.com
www.jiribenovsky.org*

NOTES

¹ See Williamson (1994), p. 151-164, for a detailed exposition of additional problems with logic and semantics that supervaluationism carries with it.

² Epistemicists' various answers to these worries include Williamson (1994), Sorensen (2001), or Weatherson (2003) – in what follows I shall present an alternative. Romerales (2004) follows, for different reasons related to the problem of higher-order vagueness, a similar strategy where there are indeterminate statements but the indeterminacy is sharply divided.

³ In this respect, the view under consideration here differs importantly from Williamson's epistemicist view which claims that our ignorance is irremediable.

⁴ Braun and Sider (2007) also use a notion of "approximate truth" in their treatment of vagueness, but it is different from the one I am appealing to here.

⁵ What about higher-order vagueness? It could be unsurprisingly an important objection here, since it is standardly taken to be a problem for supervaluationism. The problem comes from the use that both these views make of the Everest-candidates: there simply is no determinate set of candidates, for exactly as there are many candidates for being Everest, there are many candidates for being the frontier where the vague border-area of Everest stops. In short, it is indeterminate where the candidates for being Everest start and stop to be. And so, the objection goes here, how can we say that Everest_s supervenes on them, if it is indeterminate which and how many of them there are? The reply (also used by the supervaluationist) is straightforward: just resolve the problem of higher-order vagueness in the same way the problem of first-order vagueness is resolved. If you have a recipe that works in one case, then it will also work in the other case – and if the objector challenges you again with third-, fourth-, fifth-, and n-th-order vagueness, just keep using the same reply. At any level the reply will be adequate and coherent (provided it succeeds at the first level, of course). This strategy accommodates well the idea that we only can know *approximately* what Everest is: with each level of indeterminacy, we get closer and closer to Everest (that is, Everest_s), making our approximations as precise as possible for us. (See also Braun and Sider (2007), §1.4, for this strategy).

⁶ Note that supervaluationism faces a similar worry, as Tye (1990), p. 542, remarks: "[For the supervaluationist], precisification is conceived of as a process of selection. One precise object is selected from the set of precise objects associated with the given term. Now what determines membership in the relevant set? In the case of 'Everest' for example, it cannot be said that the appropriate set consists of those various objects that result from *making the reference* of 'Everest' completely precise. For this assumes that there is something imprecise that 'Everest' refers to *prior* to precisification. Nor can it be said that the relevant set is the one that consists of all those objects that 'Everest' can refer to *after* it has been made completely precise. For the process of precisification is now opaque: it involves selecting an object from a set comprised of those objects that can be selected via precisification. Unfortunately, I see no other ways of delineating the relevant sets. It seems to me, then, that supervaluationism does not successfully eschew vague objects."

⁷ I don't have in mind only direct uses of the word "Everest". I also have in mind to take into account *indirect* uses of it. For example, I never actually say, when I am in London, that I am not on Everest. But I do say that I am in Europe, and some-

times I do say that Everest is in the Himalayas, and I do say (or think) that the Himalayas are not in Europe. All this must be taken into account. Unspoken and only *thought* uses should also be taken into account.

⁸ I am grateful to Roy Sorensen for stressing this point.


⁹ For their very helpful comments on early ideas and earlier versions of this paper I would like to thank Mark Heller, Jean-Roch Lauper, and Roy Sorensen. Special thanks go to two anonymous referees of **teorema**, as well as to the Editor, whose helpful objections and comments allowed me to improve parts of the paper.

REFERENCES

- BARNES, E. (forthcoming), "Ontic Vagueness: A Guide for the Perplexed", *Noûs*.
BRAUN, D. and SIDER, T. (2007), "Vague, So Untrue", *Noûs*, vol. 41, pp. 133-156.
HELLER, M. (manuscript), "Against Epistemic Vagueness".
ROMERALES, E. (2004), "La teoría pragmática de la vaguedad. Problemas y perspectivas (The Pragmatic Theory of Vagueness. Problems and Perspectives)", *Theoria*, vol. 19, pp. 49-75.
SORENSEN, R. (2001), *Vagueness and Contradiction*, Oxford University Press.
TYE, M. (1990), "Vague Objects", *Mind*, vol. 99, pp. 535-557.
WEATHERSON, B. (2003), "Epistemicism, Parasites and Vague Names", *Australasian Journal of Philosophy*, vol. 81, pp. 276-9.
WILLIAMSON, T. (1994), *Vagueness*, London, Routledge.
WILLIAMS, R. (2008), "Ontic Vagueness and Metaphysical Indeterminacy", *Philosophy Compass*, vol 3, pp. 763-788.

WITTGENSTEIN

william child



ROUTLEDGE

