Vagueness at every order

Author

September 7, 2011

Abstract

A number of arguments purport to show that vague properties determine sharp boundaries at higher orders. That is, although we may countenance vagueness concerning the location of boundaries for vague predicates, every predicate can instead be associated with precise knowable cut-off points deriving from precision in their higher order boundaries.

I argue that this conclusion is indeed paradoxical, and identify the assumption responsible for the paradox as the Brouwerian principle B for vagueness: that if p then it's determinate that it's not determinate that not p. Other paradoxes which do not appear to turn on B turn instead on some subtle issues concerning the relation between assertion, belief and higher order vagueness.

In this paper a B-free picture of assertion, knowledge and logic is outlined which is completely free of higher order precision. A class of realistic models containing counterexamples to B and a number of weakenings of B are introduced and its logic is shown to support vagueness at every order.

It is an upshot of classical logic that if there are any small numbers at all, something I shall assume throughout this paper, then there is a last small number. It is compatible with this result that it is a vague matter which number that is. The boundary between the small and non-small isn't precise. There is a boundary, but it's *vague* where it lies, and it is the existence of precise boundaries, not vague ones, that we should be worried about.

This is, in a very schematic form, the classical response to the Sorites paradox.[1] To illustrate why precise boundaries (but not vague ones) are problematic consider the following example. There is something very bad about asserting that the total length of your childhood was 378432178928476829 nanoseconds. Vagueness prevents you from ever discovering this, and similar precise facts about the length of your childhood. If the boundary between your childhood and the rest of your life was not vague, however, there would have been no reason you couldn't have discovered the length of your childhood in nanoseconds, just as, perhaps, one could find out the number of nanoseconds in a year, and no reason to refrain from going about asserting it. Everyone, regardless of her preferred account of vagueness, must agree that the above assertion is bad and that this badness is due to its being at best vague whether your childhood lasted for this length of time.

This reasoning extends. Say that it's determinate that p just in case p and it's not vague whether p.[2] Is there a precise boundary between the determinate children, in other words, the non-borderline children, and everyone else? It seems we should not be any happier about assigning sharp numbers to the length of one's determinate childhood than to ones childhood. There is a completely analogous Sorites for 'determinate child' as there is for 'child'. To be sure, there is a last child and a last determinate child in any Sorites sequence, but it is always vague which person that last child or determinate child is. Similar comments apply to the further iterations: it's vague which the last determinately determinate child is in the sequence, and so on and so forth through the finite orders.

Indeed, there are some people who are determinately^n children for any amount of iterations, n, and some which are not. Surely it is vague where that boundary lies as well? In other words, there are some children such that it's neither vague nor higher order vague (vaguely vague or vaguely vaguely vague or ...), whether they're children, and others such that it is either vague or higher order vague whether they are children, and it's indeterminate where the boundary between the two lies. To see this note that:

> The period of my childhood during which it was neither vague nor higher order vague whether I was a child was 378432178928476829 nanoseconds in length. (1)

sounds just as terrible, and would sound terrible no matter what number one

[1] Although note that you want these claims to assuage the initial intuition that there can't be a last small number, know that you let the epistemicist off the hook too.

[2] I am thus using the phrase 'it's determinate that p' in a very neutral way which permits even an epistemic interpretation.

used. However, if neither this sentence, nor any sentence like it, were vague, what possible reason could prevent us from finding out whether or not (1) was true? If we knew (1) it would be hard to explain the inappropriateness of asserting (1) (although see [6] for a possible explanation.)

It should be noted that the only resources one needs to establish the existence of higher order vagueness is (i) an operator expression, 'it's vague whether', giving one the means to express when someone is a child without it being vague whether they're a child (that is, when someone is a determinate child, in my nomenclature), (ii) the acknowledgement that a Sorites sequence for 'x is a child' is typically also a Sorites sequence for 'x is a determinate child and (iii) that any Sorites susceptible predicate has vague instances. Since all these considerations can be made independently of ones preferred analysis of vagueness this is a problem for everybody – epistemicism, for instance, offers no respite here.[3]

Some philosophers deny the phenomenon of even second order vagueness (see, for instance, [20] and [15].) According to these theorists, the duration of your determinate childhood is a precise length. However these philosophers do not typically offer concrete hypotheses about the exact number of seconds at which it begins to be vague whether you're a child (and at which you stop being a determinate child.) If this is due to some kind of inability on their part, some explanation of this inability is required. The explanation cannot be that the boundary is vague, because by assumption the distinction between being determinately bald and being vaguely bald is a precise one. One might reasonably wonder what this new kind of obstacle to knowledge is if it is not vagueness. Equally puzzling issues arise for those who think higher order vagueness cuts out at some finite level larger than 2 (see Burgess [4].)

The subject of this paper concerns a number of arguments that purport to show that for any predicate, F, there must be a precise boundary between the things which are determinately F at every order (henceforth "determinately* F") and the rest. That being vaguely F at some order or other and not being vague at any are precise distinctions. If this argument succeeds we should expect to be seeing exact numbers associated with vague predicates all over the place. Indeed numbers that are in principle discoverable; thus one should not be surprised to hear things like 'my determinate* childhood lasted exactly 378432178928476829 nanoseconds' or 'I became determinately* bald after I lost my 1451st hair' and so on. It is tempting to sweep the infinitary version of the Sorites paradox under the carpet - to say that predicates of the form 'determinately*-F' are indeed precise but they're so esoteric we shouldn't worry about them. I think that recognising that this response involves the possibility of finding out propositions like (1) is a cost many would not be willing to pay.

Let me introduce some notation. I shall write ∇p to mean it's vague whether p, Δp to mean that p and it's not vague whether p, and $\Delta^* p$ to mean p and it's neither vague nor higher order vague whether p, which is to say p, it's not vague

[3] Eklund [9], for example, claims that his particular analysis of vagueness fairs better with respect to the problem of higher order vagueness. However, since he can make sense of the property of being bald without being vaguely bald, and he believes that Sorites susceptability involves vagueness it is hard to see how this claim stands up to this version of the problem.

whether p, it's not vague whether it's vague whether p, and so on. The problems considered can be generated without iterating into the transfinite ordinals, so by 'higher order vague' I shall just mean nth-order vague for some finite order n. Accordingly these operators can be connected in terms of an infinite conjunction as follows: $\Delta^* p \equiv \bigwedge_{n\in\omega} \Delta^n p$.

Classical logic shall be assumed throughout the paper. The claim that it is never vague whether something is determinate* is formally represented by the schema

$$\Delta\Delta^* p \vee \Delta\neg\Delta^* p \qquad (2)$$

which can be split up into two claims:

($+\Delta$) $\Delta^* p \to \Delta\Delta^* p$.

($-\Delta$) $\neg\Delta^* p \to \Delta\neg\Delta^* p$.

The former principle is introduced and argued for in Williamson [17]. Whilst some have rejected it, notably Field [12], it is not our intention to do so here. In fact an argument for ($+\Delta$) is given in §1.2. The latter principle has received less attention, although it is crucial for the problem. One way to prove ($-\Delta$) is to argue for the Brouwerian principle for Δ

(B) $p \to \Delta\neg\Delta\neg p$

Although B is the most obvious candidate (see [18], [8] footnote 11, [10]) ($-\Delta$) is also entailed by a class of weaker principles

(B^n) $p \to \Delta(q \to \phi_n)$

where $\phi_1 := \neg\Delta\neg p$; $\phi_{n+1} := \neg\Delta\neg(q \wedge \phi_n)$.[4] I shall discuss the motivations behind these principles in §2. To this list I would also like to add a principle, which I dub B^*:[5]

(B^*) $\Delta(p \to \Delta p) \to (\neg p \to \Delta\neg p)$

In this paper I shall explain how one can resist the conclusion that there is a precise boundary between the determinately* F's by rejecting these principles. I generalize some recent considerations (Mahtani [14], Dorr [7]) that show that B fails to show that B^n can fail for any n. I then consider B^* and provide tentative reasons against and in favour of it. Arguments by Wright, Graff and Zardini that do not appear to rest on B are also considered. Finally I discuss some technical issues concerning the logic of vagueness.

[4] The weakenings of B sometimes considered are the principles $\mathsf{B}^{n\prime}$: $p \to \Delta\neg\Delta^n\neg p$. These are slightly weaker (see footnote 5 for the frame conditions) but we shall see that the strengthening is more motivated in this context.

[5] In terms of Kripke frames, B^n corresponds to the condition that if x can see y, then there is a path back to x from y with at most n steps each step of which x can see, whereas B^* entails that x can see a finite path back but places no bounds on the number of steps it might take. $\mathsf{B}^{n\prime}$ is weaker that B^n stating only that there is a path with at most n-steps back, but that x needn't see any of these steps.

The structure of the paper is as follows. In §1 I outline the problem of higher order vagueness, the principles which generate it, and argue that it is B and it's weakenings that are responsible. In §2 I outline a picture of the structure of higher order vagueness in which B fails, consider the relation between vagueness, assertability and knowledge and address some other paradoxes which appear not to rely on B. Some technical questions to do with the proposed logic of vagueness are considered in the appendix.

1 The problem of higher order vagueness

Many authors have thought that precise cut-off points re-emerge once considerations involving higher order vagueness are taken into account. The following passage by Mark Sainsbury is often cited in favour of this conclusion:

> Suppose we have a finished account of a [vague] predicate, associating it with some possibly infinite number of boundaries, and some possibly infinite number of sets. Given the aims of the description, we must be able to organize the sets in the following threefold way: one of them is the set supposedly corresponding to the things of which the predicate is absolutely definitely and unimpugnably true, the things to which the predicate's application is untainted by the shadow of vagueness; one of them is the set supposedly corresponding to the things of which the predicate is absolutely definitely and unimpugnably false, the things to which the predicate's non-application is untainted by the shadow of vagueness; the union of the remaining sets would supposedly correspond to one or another kind of borderline case. So the old problem re-emerges: no sharp cut-off to the shadow of vagueness is marked in our linguistic practice, so to attribute it to the predicate is to misdescribe it. [16]

Raffman, for example, describes this reasoning as 'decisive.' However, as we have seen, it's conclusion is paradoxical. If the distinction between the people who are 'absolutely definitely and unimpugnably' bald is a precise one, we ought to be able to find out and say where it applies much like we can in principle find out if someone has less than 1000 hairs. This much is characteristic of precise distinctions.

Much turns on whether you accept classical logic or a non-classical logic. Sainsbury's conclusion that a given object either falls under a predicates realm of application, absolutely definitely and unimpugnably, or it fails to do so in some way, is equivalent to an instance of the principle of excluded middle. Perhaps there is some reason why this particular instance of excluded middle must be true, but Sainsbury's argument has done nothing to establish that. For the classical logician, however, the conclusion has no bite; what distinguishes a precise from a vague predicate is not whether it obeys the principle of excluded middle. The question is: could it be vague whether a sentence is absolutely definitely and unimpugnably true, without a shadow of vagueness?

Let us grant, for the time being, this talk of claims having 'shadow's of vagueness.' The argument could continue. Suppose it could be vague whether a sentence was absolutely definitely and unimpugnably true – true without a shadow of vagueness. The fact that it is vague whether or not the sentence has a shadow of vagueness, presumably counts as its having a shadow of vagueness. Thus it isn't absolutely definitely and unimpugnably true without a shadow of vagueness. If it's vague whether a sentence is true without a shadow of vagueness, it's not true without a shadow of vagueness.

This conclusion, however, is completely compatible with its being vague, in some cases, whether a sentence is definitely and unimpugnably true untainted by the shadow of vagueness. For according to classical treatments of vagueness it is consistent to say that something is a vague instance of a property without falling under that property. For example, suppose it is vague whether Harry is bald, and (consequently) that it's vague whether or not he's not bald. By excluded middle, either Harry is bald or he isn't. Therefore either Harry is bald and it's vague whether or not he's not bald, or Harry is not bald and it's vague whether or not he is bald (by the classical inference $p, q \vdash ((p \wedge r) \vee (q \wedge \neg r))$.) In either case we have a property, F, such that Harry vaguely instantiates it, but doesn't in fact instantiate it (F is 'not bald' in the first case and 'bald' in the latter.) Of course, there are no determinate examples of things which are vaguely F without being F, so there are no determinate examples of things such that it is vague whether they have a shadow of vagueness. If it's vague whether a claim has a shadow of vagueness, it's also second order vague (and by analogous reasoning, it's vague at all orders.) But it is also compatible with a classical treatment of vagueness that it can be determinate that there are G's without there being any determinate G's; it is no objection to this approach that we cannot find any determinate examples of the phenomenon we are interested in. To conclude that every predicate is precise at higher orders requires further consideration.

1.1 The problem of higher order vagueness

There is much in Sainsbury's argument which is left unexplained. What do the adjectives 'absolutely', 'definitely' and 'unimpugnably' add? What does he mean by truth without a 'shadow of vagueness.' In this section I shall interpret having a 'shadow of vagueness' simply as having vagueness at some order, and I shall use this to develop a more specific version of this argument.

Imagine that we are talking about the natural numbers less than 100 and we want to know which ones are small. Obviously, there's no sharp boundary between small and non-small numbers. So there will be the numbers which are definitely small, the numbers which are definitely not small, and the borderline cases in between. There's also no sharp boundary between the definitely small numbers and the rest either: there are numbers for which it's vague whether they're definitely small or borderline small. To put it another way, there are numbers which are definitely definitely small, and those which are definitely not definitely small, but there's a range of borderline cases between the two in this case as well. Similar points apply to the boundary between the definitely definitely small numbers and the numbers which are not definitely definitely small, and so on and so forth.

One can see that the set of small numbers, the definitely small numbers, the definitely definitely small numbers, and so on, gradually shrinks as the number of 'definitely's' increases; after all, being definitely F is generally a more stringent condition than being F. But this set can't shrink forever! The first set in this sequence clearly starts off with less than 10,000 members, so in this most generous case it can shrink at most 10,000 times before it becomes empty. For some N being definitely^N small is the same as being definitely^m small for any $m \geq N$.

Does this result mean there is some kind of sharp boundary between the determinately^N small numbers and the rest? Fortunately this doesn't follow from what we have said so far. We may, and indeed must, accept that there is a set of numbers which are determinately^n small for every n, and a largest such set, but we may also maintain that *it's vague which set that is.* The starting point of this shrinking process - the set of small numbers - was vague, so there is no reason to think that it won't be vague which set you end up with after the shrinking process is complete. In keeping with the denial of sharp boundaries, we may still hold that it's vague which number is the last determinately^N small number. The most we can conclude is that *if* a number is determinately^N small it is determinately^k small at all orders k.

In the following sections I shall be considering various proposals that supplement this argument to show there must be sharp boundaries. First let us make this a little bit more formal. The toy argument above made an essential appeal to the fact that the Sorites sequence considered was discrete. Not all Sorites sequences are discrete, however, for example, smallness over the rational numbers, or redness over a spectrum. A completely general argument can be found in Williamson [17] which shows that if something is definitely^n F for any number of iterations n, then it definitely is definitely^n F for every n. We may define this strong notion of definiteness using infinitary conjunction:

- $\Delta^*\phi := \bigwedge_{n<\omega} \Delta^n\phi$

In order to prove the result Williamson assumes the principle that conjunctions distribute over determinacy

(D) $\bigwedge_{i<\omega} \Delta\phi_i \to \Delta \bigwedge_{i<\omega} \phi_i$.

The argument proceeds as follows. We firstly note the logical truth: $\vdash \bigwedge_{n<\omega} \Delta^n\phi \to \bigwedge_{n<\omega} \Delta^{n+1}\phi$, which is just an instance of conjunction elimination.[6] From (D)

[6] This may be proved from the principles C1-C3 below

C1. $\bigwedge_{i<\omega} \phi_i \to \phi_n$ for each $n < \omega$.

C2. $\bigwedge_{i<\omega}(\phi \to \psi_i) \to (\phi \to \bigwedge_{i<\omega} \psi_i)$.

C3. If $\vdash \phi_i$ for each $i < \omega$, $\vdash \bigwedge_{i<\omega} \phi$.

$\vdash \bigwedge_{n<\omega} \Delta^n\phi \to \Delta^i\phi$ for each $0 < i < \omega$ by C1. Thus $\vdash \bigwedge_{i<\omega}(\bigwedge_{n<\omega} \Delta^n\phi \to \Delta^{i+1}\phi)$ by C3. So finally $\vdash \bigwedge_{n<\omega} \Delta^n\phi \to \bigwedge_{i<\omega} \Delta^{i+1}\phi$ by C2.

we can immediately infer $\vdash \bigwedge_{n<\omega} \Delta^n\phi \to \Delta \bigwedge_{n<\omega} \Delta^n\phi$, i.e.

$(+\Delta)$ $\vdash \Delta^*\phi \to \Delta\Delta^*\phi$.

It should also be noted that none of these principles are characteristically classical, so this result extends to a number of non-classical logics.

1.2 Is a conjunction of determinate truths determinate?

The most natural place to block Williamson's argument is to deny (D). This is Field's strategy in [12]. But is this denial at all plausible? I claim it isn't if we accept the following principle concerning vagueness:

> If each constituent of a sentence is precise then the sentence itself is precise. (3)

Suppose for reductio that $\bigwedge_{i\leq\omega} \Delta\phi_i$ but that $\neg\Delta \bigwedge_{i\leq\omega} \phi_i$. Since each ϕ_i and infinitary conjunction are precise, it follows that $\bigwedge_{i\leq\omega} \phi_i$ is precise by (3), i.e. $\Delta \bigwedge_{i\leq\omega} \phi_i$ or $\Delta\neg \bigwedge_{i\leq\omega} \phi_i$. By assumption it's not true that $\Delta \bigwedge_{i\leq\omega} \phi_i$ so $\Delta\neg \bigwedge_{i\leq\omega} \phi_i$ must hold. By factivity we have ϕ_i for each i and $\neg \bigwedge_{i\leq\omega} \phi_i$ - this is a contradiction by C3.

It is interesting to note that if you accept the principle B, you can actually *prove* that a conjunction of determinate truths is determinate from the logic of conjunction alone (see appendix 3.2.) The thesis of this paper, however, is that once you have rejected B, you already have a satisfying response to the paradoxes of higher order vagueness available. It is to this principle we now turn.

1.3 Sharp boundaries from B

To show that 'determinately* small' is a sharp predicate we must show that it can never be vague whether something is determinately* small. We have so far shown that if a number is determinately* small then it's not vague whether it's determinately* small. If we could show that if a number is not determinately* small then it's not vague whether it's determinately* small, we would be able to get our conclusion from an instance of excluded middle and reasoning by cases: every number is either determinately* small or it isn't, and in either case it is not vague whether it's determinately* small. The arguments I shall consider work in favour of the completely general principle I have called $(-\Delta)$: $\neg\Delta^*p \to \Delta\neg\Delta^*p$. In conjunction with $(+\Delta)$: $\Delta^*p \to \Delta\Delta^*p$ we then have the problematic principle $\Delta\Delta^*p \vee \Delta\neg\Delta^*p$ stating that there is never vagueness concerning what is determinate*.

A simple way to close this gap would be to introduce the principle B:

$$\mathsf{B}: \neg p \to \Delta\neg\Delta p \qquad (4)$$

Here is how you prove $(-\Delta)$ from B: we have $\neg\Delta^*p \to \Delta\neg\Delta\Delta^*p$ by B. By contraposing $(+\Delta)$ we have $\neg\Delta\Delta^*p \to \neg\Delta^*p$, so we also have $\Delta\neg\Delta\Delta^*p \to \Delta\neg\Delta^*p$

by necessitation and an application of K. So by transitivity of the conditional we have $\neg\Delta^* p \to \Delta\neg\Delta^* p$.[7] Indeed, it is exactly this principle which is assumed in Williamson's discussion of these issues in [18].

The axiom B is motivated by Williamson's fixed margin models described in [17]. Williamson describes the class of Kripke frames $\mathcal{C}$ which contains frames, $\langle W, R\rangle$, for which there is some metric over W, $d(\cdot,\cdot)$, and $\alpha \in \mathbb{R}$ such that Rxy iff $d(x,y) \leq \alpha$. The motivation for this semantics is roughly the same whether we are epistemicist or some kind of supervaluationist. We should think of the members of W as precise interpretations of the language with the metric representing the degree of similarity between two interpretations. For example, suppose two interpretations, x and y, agree on how to interpret every expression, except x interprets 'small for a number less than 100' as the numbers less than 14 and y as the numbers less than 13. These two interpretations should be considered quite close by the intended metric, which is to say that $d(x,y)$ will be a relatively small number. A formula, ϕ, is determinately true according to an interpretion, x, if ϕ is true at all the interpretations similar enough to x. 'Similar enough' means the measure of their differences does not exceed α. At the interpretation x described above, '2 is small for a number less than 100' is determinately true because interpretations that disagree with x on this sentence are quite far away.[8]

It is clear that the accessibility relation of each such frame is symmetric: if the distance between x and y is less than α then, the distance between y and x is less than α. It is a standard fact that B is validated in all and only symmetric frames, so for any frame $\langle W, R\rangle$ in $\mathcal{C}$ the axiom B is valid.

1.4 Mahtani on failures of B

In [14] Mahtani argues that since the term 'determinately' is itself vague, its interpretation ought to vary from point to point and that Williamson's models fail to capture this fact. The interpretation of 'determinately' technically does vary from interpretation to interpretation on Williamson's semantics - being determinate at x depends essentially on which interpretations are closest to x. However they do not vary in all the salient respects. In particular, all interpretations agree about how close an interpretation has to be to be 'close enough' in the relevant sense; α is a fixed quantity throughout the model. In Mahtani's terminology the 'accessibility range' does not vary, when it should.

If each point, x, has its own accessibility range, $f(x)$, symmetry is no longer guaranteed. The distance between x and y may be less than $f(x)$ but not less than $f(y)$ (see figure 1.)

[7] It is interesting to note that B is not so central for the non-classical theorist. The problematic principle is the rule: if $\vdash \phi \to \Delta\phi$ then $\vdash \phi \vee \neg\phi$. For example both Łukasiewicz and Field's recent logic have this principle. Once one has this principle and $(+\Delta)$ all the problematic classical theorems concerning Δ^* statements are provable. Things would be different, however, if we adopted the weaker rule: if $\vdash \phi \to \Delta\phi$ and $\vdash \neg\phi \to \Delta\neg\phi$ then $\vdash \phi \vee \neg\phi$.

[8] The closest interpretation to x which disagrees interprets 'is small for a number less than 100' as the numbers less than 2, whereas x interprets that as the numbers less than 14. In this context this constitutes a fairly big difference.

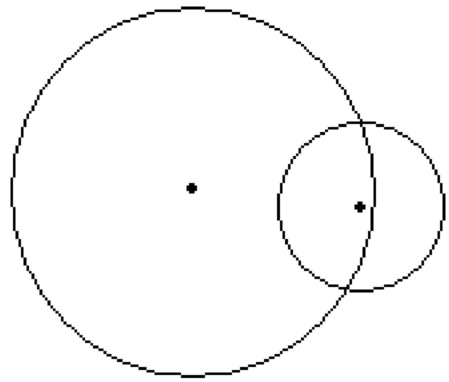

Figure 1: Mahtani's counterexample to B

1.5 Sharp boundaries from other principles

Our initial concern was whether it must always be a precise matter whether something is determinate*. We noted that B was one way, but not the only way, to close the gap between Williamson's argument for (1) and this claim.

Unfortunately there is an infinite chain of weaker principles, all stated in the finitary language, that also prove that 'determinate*' is precise.

$$\mathsf{B}^n\text{: } p \to \Delta(q \to \phi_n) \tag{5}$$

where $\phi_1 := \neg\Delta\neg p$; $\phi_{n+1} := \neg\Delta\neg(q \wedge \phi_n)$. And the principle

$$\mathsf{B}^*\text{: } \Delta(p \to \Delta p) \to (\neg p \to \Delta\neg p) \tag{6}$$

For example, adding B^* to the infinitary language allows us to prove the problematic $\neg\Delta^*\phi \to \Delta\neg\Delta^*\phi$. By applying necessitation to $(+\Delta)$, which we are assuming at this point, we have: $\vdash \Delta(\Delta^*\phi \to \Delta\Delta^*\phi)$. However, $\Delta(\Delta^*\phi \to \Delta\Delta^*\phi) \to (\neg\Delta^*\phi \to \Delta\neg\Delta^*\phi)$, is an instance of B^* in the infinitary language, so we can imediately infer $(-\Delta)$, i.e. $\neg\Delta^*\phi \to \Delta\neg\Delta^*\phi$, by modus ponens, as required.

In terms of frame conditions, B^n characterises the the property that if Rxy then there are n points, $z_1, \ldots, z_n$ such that Ryz_n, Rz_nz_{n-1}, …, Rz_2z_1 and Rxz_i for each i - i.e. if x sees y then you can get back from y to x in n steps which x can see. B^* characterises the the property that if Rxy then for *some* n, there are $z_1, \ldots, z_n$ such that Ryz_n, Rz_nz_{n-1}, …, Rz_2z_1 and Rxz_i for each i - i.e. if x sees y then you can get back from y to x in finitely many steps which x can see. Call this latter property the 'backtrack' property. In the lattice of modal logics, KTB^* is the infimum of $\{\mathsf{KTB}^n \mid n \in \omega\}$.

As stated, any one of these axioms is sufficient to close the gap between (1) and the existence of sharp higher order cut-off points. I shall show that each frame validating KTB^n or KTB^* also validates $\neg\Delta^* p \to \Delta\neg\Delta^* p$ and hence $\Delta\Delta^* p \vee \Delta\neg\Delta^* p$. Suppose the frame $\mathcal{F} := \langle W, R\rangle$ validates KTB^n or KTB^*. For any model based on $\mathcal{F}$, if $\neg\Delta^* p$ is true at x and Rxy then (a) for some n you can get from x to a $\neg p$ world in n steps. (b) for some m you can get back to x from y in m steps. Thus you may get from y to a $\neg p$ world in $n+m$ steps, so $y \nVdash \Delta^{n+m} p$ and thus $y \Vdash \neg\Delta^* p$. But y was an arbitrary point accessible from x, thus $x \Vdash \Delta\neg\Delta^* p$. So $x \Vdash \neg\Delta^* p \to \Delta\neg\Delta^* p$ for every x.

It should be clear by now that to properly demonstrate that we can consistently deny higher order precision we will need a more principled way of generating counterexamples. This is what I shall attempt to do in the next section.

2 A B-free solution

To deny sharp cut-off points at all levels we must reject the principles B^n and B^*. The rest of this paper is an evaluation of the prospects of this proposal. The general strategy is to find models in which the claims 'there are borderline cases of being determinately* (determinately^n) small' ($\exists x \nabla\Delta^* Sx$, $\exists x \nabla\Delta^n Sx$) come out true. These models serve the purpose of showing that no contradiction can be derived from these assertions and plausible background assumptions. I attempt also to construct more realistic models, taking Williamson's fixed margin models as a starting point, to show that these principles are compatible with more realistic structural assumptions.

Below is a model which demonstrates that it is at least possible to deny sharp cut-off points at every level. Each node represents an interpretation of the language, with the number at each node representing the greatest number which satisfies 'small for a number less than 100' according to that interpretation. The truth value of each atomic statement of the form 'n is small' at a node is thus determined by whether number n is less than or equal to the number at that node. The truth values of extensional combinations of formulae are calculated as usual relative to a node, and a claim of the form '$\Delta\phi$' is evaluated true at a node x iff ϕ is true at every node accessible via an arrow from x.

No point can see a point which differs from it by more than one. The converse fails, however, since interpretations may differ radically in the interpretation of other expressions and might thus be inaccessible to one another. It is tacitly assumed that every point sees itself.

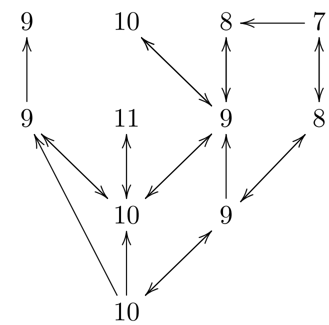

Remember that 'the last small number' according to a point is the number written beside it, so 'the last determinately* small number' at a point is the smallest number you can get to from that point by following the arrows. Note that the bottom node can see a node where the last determinately* small number is 7: follow the arrow to the right. But it can also see two nodes where the last determinately* small number is 8: follow the arrow up or left. Thus it is vague, at this point, whether the last determinately* small number is 7 or 8. Indeed, it's vague whether it's determinate* that 8 is small.

Of course we tried to make this model look realistic by having several points (we could have gotten away with two) and making sure that adjacent points didn't disagree substantially over the interpretation of 'small for a number less than 100'. However, it would be nice to have a general class of models that includes such models as a special case, but are also constrained by facts about vagueness in the same way Williamson's semantics was. In fact we can modify Williamson's fixed margin models in just the way Mahtani suggests to allow for variation in accessibility range. We must also make sure that close points don't interpret 'determinately' drastically differently, i.e. we must make sure close points have similar accessibility ranges. This motivates the following definition.

Definition 2.0.1. *A v-frame is a triple* $\langle W, d(\cdot,\cdot), f(\cdot)\rangle$ *where* $\langle W, d\rangle$ *is a metric space, and* $f : W \to \mathbb{R}$ *obeys the following:*

(A) $\forall w, v \in W, |f(w) - f(v)| \leq d(w, v)$

A formula of propositional modal logic is valid on a v-frame $\langle W, d, f\rangle$ iff it is valid on the Kripke frame $\langle W, R\rangle$ where Rxy iff $d(x, y) \leq f(x)$. We shall talk about a v-frame and it's associated Kripke frame interchangeably from now on.

The elements of W are to be thought of as interpretations or precisifications of the language, with the metric d representing how close they are to one another. A formula is determinately true at an interpretation if it is true at all nearby interpretations. What counts as nearby the interpretation w is determined by $f(w)$: v is nearby w when the distance between them according to d is less than $f(w)$. Note that what counts as 'nearby' depends on the precisification - the constraint (A) says, roughly, that the closer two interpretations are, the less they can differ over their interpretation of 'nearby'.

A useful fact is that there is a natural way to assign a metric over a (generated) Kripke frame. Simply assign a length to each arrow and define the distance between x and y to be the length, ignoring the direction of the arrows, of the shortest path between x and y. With a bit of fiddling one can show that the model above is generated by a v-frame in this way. The fact that one can refute B^n and B^* over v-frames follows also from the more general fact that the logic of v-frames is KT (see the appendix.)

2.1 Realistic frames

Let us consider a toy propositional language whose only atomic sentences are English sentences of the form: 'a is red', where a ranges over names for colours in a fixed colour spectrum. It is natural to suppose the interpretations of this language are completely specified by the cutoff point for 'red' along this spectrum of colours. Each colour can be represented by a real number, and the distance between two interpretations can be modelled as the difference between the two numbers representing the cutoff points for those interpretations. Thus the metric of our v-frame is the standard notion of distance on $\mathbb{R}$.

This suggests a very natural class of v-frames for modelling vague languages: those based on Euclidean space, $\mathbb{R}^n$, where the dimension n is the number of vague predicates in the language. The good news about these v-frames is that each of the problematic principles B^n are refutable. To refute B^n we consider the standard metric over $\mathbb{R}$. Let $\epsilon = \frac{1}{2^{n+1}}$. Stipulate that $f(0) = 1$, that $f(x) = f(-x) = \epsilon$ for $x \in \mathbb{R}\backslash[-1+\epsilon, 1-\epsilon]$ and $f(x) = 1 - x$ for $x \in (0, 1-\epsilon]$ and $x - 1$ for $x \in [\epsilon - 1, 0)$. It is easy to check this satisfies condition (A) and is a v-frame. Now 0 can see 1, yet the shortest path back from 1 takes $n+1$ steps, thus B^n does not hold. Note that in this model there are no points which can only see themselves, i.e. no points, x, such that $f(x) = 0$.

The counterexamples traded on the idea that for any n one can find a v-frame based on an ϵ small enough to ensure that the longest path from 1 to 0 is longer than n. There is, however, no single model which refutes all the B^n simultaneously. Indeed it is not hard to show that v-frames based on $\mathbb{R}^n$ where no points have a 0 accessibility range has the backtrack property: if x sees y, then there is a finite path $z_0, \ldots, z_n$ such that y sees z_0, z_0 sees z_1, ..., z_n sees x and x sees each z_i. Thus it follows that B^* holds in these frames (see figure 2.) Intuitively, the closer a point is to x, the closer in diameter it's accessibility range must be to x's thus one always can find a path leading back to x:

For our purposes this result would be devastating. It would entail, for example, that being determinately* red had completely precise boundaries, and this brings along with it all the problematic consequences already mentioned.

An obvious place to resist the result is to deny that the accessibility range of any point must be non-zero. Relaxing this constraint is independently motivated. Suppose we are working with the toy language described above, and the ordering of the relevant colour spectrum is not only dense but complete in the sense of containing a limit for any converging sequence of points (thus,

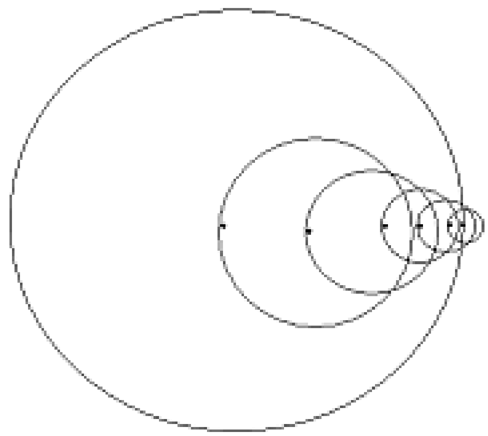

Figure 2: The validity of B^*: x sees a far out point on the edge, yet there is a finite path back to x.

for example, it's structure is like that of $\mathbb{R}$, but not of $\mathbb{Q}$.)[9] If one made the assumption that $f(x) > 0$ for any interpretation x, it follows that *no* colours in the spectrum are determinately* red. It is sufficient to show that any two interpretations, represented as real numbers, x and y, can be connected by a path. Without loss of generality we may assume that $x < y$. Let $(a_n)_{n\in\omega}$ denote the sequence $x, x + f(x), x + f(x) + f(x + f(x)), \ldots$, i.e. $a_0 := x$ and $a_{n+1} = a_n + f(a_n)$. If $a_n < y$ for each n then $(a_n)_{n\in\omega}$ must clearly converge as it is a bounded monotonic sequence. Let a_∞ be the point it converges to. I claim that $f(a_\infty) = 0$, contradicting the assumption that $f(x) > 0$ for any x. By condition (A) on v-frames we know that $|f(a_\infty) - f(a_n)| \leq a_\infty - a_n (= d(a_\infty, a_n))$ for each n. However, since the right hand side converges to 0 as n increases, and $f(a_n)$ converges to 0, it follows that $f(a_\infty) = 0$.

Once one moves away from the simple toy example, v-frames based on $\mathbb{R}^n$ become implausible for other reasons. For example, suppose now we are considering a language in which the only atomic sentences are of the form 'a is red' and 'a is orange' for a a colour in a fixed spectrum. Modelling this language using $\mathbb{R}^2$ would be overly simplistic because the interpretation of 'red' and 'orange' are not independent. Any assignment of cutoff points that allowed 'red' and 'orange' to overlap should be intuitively very far away from the intended interpretation. Thus an interpretation that says that the red colours end at the colour representend by 10 and orange starts at the colour represented by 9 should be very far away from the sensible interpretation that says red ends

[9]If one were to object that some fact about colours prevents the existence of such a spectrum, we could modify the example to be about the vague predicate 'small' as applied to real numbers.

at 9 and orange starts at 10. However their distance according to the standard metric on $\mathbb{R}^2$ is relatively small: $\sqrt{2} = \sqrt{(10-9)^2 + (9-10)^2}$.

A final worry in this ballpark is that even if $\mathbb{R}^n$ v-frames aren't suitable for modelling vagueness, the correct models might still have enough $\mathbb{R}^n$-like properties to guarantee that the backtrack properties hold. Let me finish by considering two such properties we might quite plausibly expect to hold in any realistic model:

- Density: for any x and y there is a z such that $d(x, z), d(y, z) < d(x, y)$.
- Closeness: Whenever $d(x, y) \leq f(x)$ there is a z such that $d(y, z) \leq f(y)$ and $d(x, z) < d(x, y)$.

However, neither of these principles, even in tandem, ensure that the relevant v-frame has the backtrack property. A simple example would be to let $W := [0, 1) \cup (2, 3]$, $d(x, y) = |x - y|$, $f(x) = 1.5$ for $x \in [0, 1)$ and $f(x) = 1$ for $x \in (2, 3]$. Anything in the range $(\frac{1}{2}, 1)$ can see points in $(2, 3]$, but there is no path from a point in $(2, 3]$ to a point in $[0, 1)$. Furthermore this v-frame is both dense and close. These examples are *gappy*: for a given point x, there may be a range of real numbers $[\alpha, \beta]$, such that the distance between x an another point is never in $[\alpha, \beta]$.

2.2 Precise ranges

When considering the simple model presented in this section one might object that it is no better to say that determinately 7 or 8 is the last determinately* small number (although it's indeterminate which) than to say that determinately 7 is the last determinately* small number. Indeed, it seems just as bad to say that the location of this cut-off point is vaguely located over a precise range as it is to say that it is precisely located somewhere.

One might question the intuition that this is just as bad. After all, if we consider the vague predicate 'Small for a number less than 10' it seems fair enough to say that, determinately, either 3 or 4 is the last small number less that 10, although it's indeterminate which. 'small for a number less than 3' seems to be precise.

However there is no need to challenge the intuition. For according to that model there is no precise range in which the last determinately* small number vaguely falls. At the leftmost and middle point the bottom node can see it's vague whether the last determinately* small number is 8 or 9, and at the rightmost point it's vague whether it's 7 or 8. Indeed this is not just a peculiarity of our model. The contrapositive of $(+\Delta)$ says $\neg\Delta\Delta^* p \to \neg\Delta^* p$ so can show fairly easily that $\nabla\Delta^* p \to \nabla\nabla\Delta^* p$.[10] In other words, if it's vague whether p is determinate* then it's vaguely vague.

[10]Proof: let q be $\Delta^* p$, so that we have $\neg\Delta q \to \neg q$. So $\nabla q \to (\neg q \wedge \nabla q)$ by definitions and positional logic. Since we can prove the consequent is not determinate, we can prove in K that the antecedent is not determinate, i.e. we can prove $\neg\Delta\nabla q$, so $\nabla q \to \neg\Delta\nabla q$. Since $\nabla q \to \neg\Delta\neg\nabla q$ we have $\nabla q \to \nabla\nabla q$. Indeed, given the equivalence between q and $\Delta^n q$, it is possible to show that $\nabla q \to \nabla\Delta^n q$ for any n whatsoever.

How much of a limitation does this place on our solution? Must we retract all our claims about there being no sharp cut-off points for 'determinately* small' - must we retract them not because they are false, but because they are vague? Evidently we must retract specific claims of the form 'it is vague whether p is determinate*' since they are all at best vague. I claim this is an advantage of the theory, even, since we can avoid the objection that we must have a precise range in which the last determinately* small number falls.

What about the crucial claim that it is vague where the last determinately* small number lies? This claim still stands, and is determinate. In our model not only is it vague, at the bottom node, where the last determinate* number lies, but it's determinately vague where the last determinate* small number lies. This should not be puzzling for the classicist - it is standard to allow it to be determinate that there are F's whilst denying the existence of a determinate F. It is no different for the complex property of being vaguely determinately* small.

2.3 The Forced March Sorites and Assertion

One of the more puzzling issues relating to vagueness and higher order vagueness is the so-called 'forced Sorites march.' One is to imagine that you are to be presented with the elements of a Sorites sequence for F in succession, and in each case you are required to say to the best of your ability whether the element is F or not. The puzzle is that there must surely be a first element at which you stop saying 'yes, it's F' and switch to doing something else. Perhaps that is saying 'no, it's not F', or saying 'I don't know' or perhaps it is not saying anything at all. The point is, whatever one does, it seems one is committed to a sharp boundary. (For simplicity I shall confine attention to the responses: 'yes', 'no', 'I don't know', 'it's indeterminate' and saying nothing at all. Other responses are available: 'I'm not sure whether I know or not', 'I don't know whether it's determinate', 'it's indeterminate whether it's indeterminate', and so on, but since these iterations get less and less relevant to the question you're being asked to answer, silence is often a better response even if these answers are known.)

The notion of commitment here is a pragmatic one: which propositions are assertable, i.e. which answers are appropriate in a forced march, depends on the truth of the commitments of an assertion of that proposition. An assertion that p commits you to q iff, had you been as knowledgeable as possible, your assertion would have been appropriate if and only if q.[11] The view I defend states that the strongest proposition an assertion that p commits you to is the proposition that it is determinate that p. Provided you're as knowledgeable as you can be about the relevant facts, p is assertable (ignoring other pragmatic factors) when p is determinate. Thus, for example, saying 'yes, a is F' to a

[11] If you are not as knowledgeable as you could be then we have only the principle that your assertion is appropriate *only if* its commitments are true. It is only if you have the relevant background knowledge that the truth of the assertions commitments guarantees the appropriateness of the assertion as well.

given question commits you to a being determinately F – you should not say this if you think it's vague whether a is F. If you say 'it's vague' then you are committed to it's being determinate that it's vague whether a is F, and so on. Not saying anything is importantly different from saying 'I don't know', for the latter commits you to the determinacy of your not knowing, which (I shall argue) entails that you have the ability to know that you don't know. If you don't know whether you know something you are better of staying quiet than saying 'I don't know.'

My case revolves around the principles like the following

> If it's not vague whether Harry is bald, i.e. if it's a precise matter whether Harry is bald, then one can in principle know whether Harry is bald. (7)

Why might someone object to this principle? Of course, the *general* principle that all precise propositions have knowable truth values might fail for reasons not relating to vagueness at all. But I take it that if the above instance fails at all it fails for reasons relating to the vagueness of 'bald'.

I espouse the principle. According to an opposing view it might be impossible to find out whether Harry is bald, even if it is a determinate matter whether he's bald. According to classical treatments of vagueness it is possible for a proposition to be both precise, and for it to be vague whether the proposition is precise. In these cases, the opposing view claims, it is not possible to know whether p, *even though* p is precise. In these cases the obstacle to our knowing is not vagueness (the proposition in question is precise) but second order vagueness. It is very natural, on this view, to also count higher orders of vagueness as sources of ignorance in the same way.[12]

The view that vagueness and only vagueness is responsible for ignorance gives a more intuitive explanation of the phenomenon. The easiest way to explain to a non-philosopher the philosophical notion of vagueness is by saying it's whatever prevents us from knowing whether a given person is bald when we know all the relevant facts about hair number and distribution. Clearly the kinds of reasons we can't know whether a given person is bald or not, even if we know all the facts about hair number and distribution, forms *some* sort of natural kind. The simplest theory simply identifies vagueness with this obstacle to knowledge.[13] As we proceed, we shall see that the simple view does a much better job at explaining the puzzles of higher order vagueness.

The converse of (13) seems to be relatively uncontroversial, so given what I have said so far it is reasonable to assume that both (13) and it's converse are determinate – in other words the following principle is determinate: one can in principle know whether Harry is bald if and only if it's a precise matter

[12]An explicit example of this view can be found in Fine [13].

[13]One might prefer to simply introduce a separate name for whatever this obstacle to knowledge is: schmagueness. Arguments analogous to those in this paper would demonstrate that schmagueness iterates non-trivially, and strictly analogous puzzles for higher order schmagueness arise as for vagueness.

whether he is bald.[14] It is a consequence of this, in the standard logic of determinacy (which includes the K principle), that if it's vague whether Harry is determinately bald, it's vague whether we can know that he's bald. Assuming that the respondent knows as much as she can about Harry's head, it follows that it's vague whether she in fact knows that he's bald, and it thus also presumably vague whether it's permissible to assert that Harry is bald in such circumstances. None of this should be intrinsically surprising given the inherent vagueness of the relevant notions of possibility, knowledge, permission and assertion.[15] In summary: in a forced march, if you know as much as you can about the situation then 'it's determinate that p', 'you know that p' and 'it's assertable that p' are all coextensive operators.

The companion to this view about knowledge is the view that the strongest proposition an assertion that p commits you to is the proposition that p is determinate. If one is as knowledgeable as possible, then one knows p iff p is determinate. Since, ignoring other pragmatic factors, an assertion is appropriate iff it's knowledgeable, it follows that an assertion that p commits you to no more than the claim that p is determinate in a forced march. In contrast, the companion to the opposing view about knowledge has it that the strongest proposition an assertion that p commits you to is the proposition that p is determinate*.

Let us apply this theory of commitment to the forced march Sorites. We may consider the possible responses in turn. Certainly saying 'yes' up to a certain point, and then saying 'no' commits one to sharp boundaries, for that is to commit one to the elements being determinate cases up to a certain point, and determinate non-cases from there on. This is just what it is to say that F is sharp. To respond by saying 'yes' up to a certain point, a_n, and then continue by saying 'it's indeterminate' or 'I don't know', is slightly more complicated. This commits one to the determinacy of $Fa_1, \ldots Fa_n$. Asserting that it's indeterminate that each of $a_{n+1} \ldots a_m$ is F commits one to the determinacy of $\nabla Fa_{n+1} \ldots \nabla Fa_m$. Thus we are committed to the following: ΔFa_n and $\Delta\neg\Delta Fa_{n+1}$. This is not a commitment to sharp boundaries, but in a standard Sorites (without gaps) one would not expect there to be a determinate F adjacent to a determinately not determinate F. It is thus plausible to assume that this response pattern commits you to some falsehoods; this this pattern of assertions is inappropriate. Finally, if we assume the respondent knows as much as she can about the relevant background precise facts, then her knowledge is coextensive with what's determinate, so if she answers 'I don't know', we may reason as above.

On the other hand, saying 'yes' up to a certain point, and then saying nothing

[14]I am using 'can' here in a way that guarantees that you cannot know that p if p is false. 'can' cannot simply mean metaphysical possibility here because many false propositions are true, and even known, in some metaphysically possible world.

[15]I take it that the same Sorites sequence that shows that 'x is bald' is vague, shows that 'it's possible to know that x is bald' and 'it's appropriate to assert that x is bald' are vague. I am further claiming that the vague instances of the latter two predicates are precisely the second order vague instances of the former, and the determinate instances of the latter are the determinately determinate instances of the former.

for a stretch is a different matter altogether. This response pattern does not commit one to sharpness of any kind. Not saying anything is not the same as asserting that you don't know Fa_n, since the latter is not appropriate when you do not know that you don't know that Fa_n. Saying nothing does not commit you to any claim about the vagueness or nth order vagueness of Fa_n. It is completely compatible with this response pattern that every predicate of the form $\Delta^n Fx$ is vague. In other words, asserting that the cases are F up to a certain point, remaining silent for a while, asserting the cases are indeterminate, sliding back into silence for a bit, and then asserting the cases are not F from then on, does not commit you to any thing inconsistent with the thesis that $\Delta^n Fx$ is vague for each n.

To demonstrate this, suppose that $Fa_1 \ldots Fa_n$ are determinate, and that $Fa_{n+1} \ldots Fa_m$ aren't. Since there is second order vagueness, it is vague which number n is in our example, but we may be certain there is *some* such n by classical logic. If I happen to say 'yes' from cases a_1 to a_n and remain silent for cases $a_{n+1} \ldots a_m$ I have (a) asserted correctly, in the sense that I have asserted p when p is determinate and (b) have committed my self to nothing incompatible with vagueness at all orders. Furthermore, despite the fact that a_n is the last determinate F, it is presumably vague that a_n is the last determinate F, so my assertions, despite being correct, fail to be determinately correct. If a perfect asserter is someone who asserts p just in case it's determinate that p, there can be perfect asserters but it will always be at best vague whether you're a perfect asserter (provided we assume that it is always a determinate matter whether you have asserted p or not.[16])

It is worth remarking that if you knew you were a perfect asserter, i.e. if you knew that you asserted p just in case it's determinate that p, you would be able to infer from your having not asserted Fa_{n+1} that Fa_{n+1} was not determinate. Since in a typical forced march Sorites, it is at best vague whether you are a perfect asserter, such knowledge would not be available to you. So while there always is a correct response to the forced march Sorites, second order vagueness makes it impossible to know if you've made the correct responses.

2.4 Other paradoxes

There are a number of other paradoxes of higher order vagueness in the literature which do not rely on the principle B which I shall turn to now. They are both variations on an argument originally due to Wright [19]. We shall see that both these arguments rely on the view about commitment, assertion and knowledge which we have rejected in the previous section. Once the alternative is taken into account it is seen that it is compatible with these results that the thesis that there is vagueness at all orders is both assertable and known.

Let me begin with an argument due to Delia Graff Fara [11]. Fara's argument shows that natural principles concerning higher order vagueness, so called 'gap

[16] This assumption may not be unassailable. For example, I might falter or hesitate as I say 'yes' in such a way as to make it vague whether I actually committed myself to the Fness of the case in question. This might be one way to be a determinate perfect asserter.

principles' $\Delta\Delta^n Fa_k \to \neg\Delta\neg\Delta^n Fa_{k+1}$, with the rule of proof Δ-intro:

If $\Gamma \vdash \phi$ then $\Gamma \vdash \Delta\phi$

lead to a contradiction. My focus here will be the rule of proof of Δ-introduction.

There has been some debate concerning whether a supervaluationist should accept classical logic, where 'classical logic' is construed broadly to include classical rules of inference and rules of proof. The rule Fara appeals to is incompatible with certain classical rules of proof. For example one could not apply conditional proof to $p \vdash \Delta p$ (which can be obtained by Δ intro on $p \vdash p$) to obtain $\vdash p \to \Delta p$, since this would imply, as a matter of logic, that everything is precise.

There are some tricky questions in this area concerning the nature of logic, whether Δ is a logical constant, and on the normative impact both notions of consequence have. Since Fara's conclusion is essentially normative - that it is in some sense incoherent to accept vagueness at all orders - it would be nice to talk directly about the normative conclusions without the detour through 'consequence' talk which can be quite obscure in these contexts anyway. Let me introduce two notions of a 'good inference' from $p_1 \ldots p_n$ to q.

1. $Cr(q) = 1$ if $Cr(p_i) = 1$ for each $i \leq n$ and $Cr \in \mathsf{E}$.
2. $\Sigma_{i \leq n}(1 - Cr(p_i)) \leq 1 - Cr(q)$ for every $Cr \in \mathsf{E}$.

Here E is a set of credence functions that you would be justified in having in some possible epistemic situation. One notion governs what can be inferred given what you are already fully justified in believing, whereas the other constrains your beliefs when you are less than certain in the premisses. Something like these two notions are sometimes characterised in terms of global and local consequence, corresponding to the definiteness of the premisses strictly implying the definiteness of the conclusion (i.e. 'preservation of supertruth') and the premisses simply strictly implying their conclusion (i.e. 'preservation of disquotional truth'.) In formal terms that is $\Box(\Delta p_1 \wedge \ldots \wedge \Delta p_n \to \Delta q)$ and $\Box(p_1 \wedge \ldots \wedge p_n \to q)$ where $\Box$ represents some suitable notion of logical necessity.

Since one could never find oneself in an epistemic situation which *fully* supported $p \wedge \neg\Delta p$, but one could quite easily have evidence for $\neg p \vee \Delta p$ the former notion of good inference invalidates reductio: we have $p \wedge \neg\Delta p \vdash$ but not $\vdash \neg(p \wedge \neg\Delta p)$.

I am happy to engage in either talk provided it is clear what one means and one is careful which normative conclusions one draws. However, neither notion permits the inference from p to Δp. Crucially the first notion, 1., does not permit this inference. Observe first that this rule does not preserve determinate truth, for example if p is determinate but not determinately determinate, the premise of $p \vdash \Delta p$ is determinate and it's conclusion isn't. Similarly, since p is precise (although not determinately so), one could in principle have evidence which justifies certainty in p, yet be uncertain in Δp due to the vagueness in

Δp. To be sure this counterexample relied on higher order vagueness, but this is clearly not the place to beg *that* question.

Note that the consequence relation that instead preserves determinacy* does appear to validate Δ-intro. For example, according to this relation $(+\Delta)$ guarantees that $p \vdash \Delta p$. As I have argued in the previous section, if your beliefs are 'inconsistent' according to this consequence relation you are not necessarily being incoherent by committing yourself to a contradiction. The most any belief or assertion commits you to is the determinacy, not the determinacy*, of what is believed or asserted.

In this vein Zardini presents an argument that does not rely on Δ-intro [21]. His argument operates directly with the notion of determinacy*. His argument shows that if we assume the determinacy* of (a) the vagueness of $\Delta^n Fx$ for each n, (b) the Fness of a_0 and (c) the non-Fness of $a_{1,000,000}$ we can derive a contradiction. More formally he assumes the following: $\Delta^*\exists x\nabla\Delta^n Fx$, $\Delta^* Fa_0$, $\Delta^*\neg Fa_{1,000,000}$.

What is surprising about Zardini's argument is that although I have throughout been arguing for the vagueness of predicates of the form $\Delta^n Fx$, i.e. for the claim $\exists x\nabla\Delta^n Fx$, I cannot maintain that this claim is determinate*. This is interesting as it is a concrete example of something I would count as a permissible assertion which is not determinate*.

To check that these assertions, including $\exists x\nabla\Delta^n Fx$ and $\exists x\nabla\Delta^* Fx$, do not commit us to any contradictions we need to show that not only are there models in which they are all true, but that there are models in which they are all determinately true. In fact, one can show for any $n \in \omega$ that there is a model in which these claims are determinate^n true.

Let's start with an example in which $\exists x\nabla\Delta^* Fx$ is simply true. Note also that all these examples apply also to $\exists x\nabla\Delta^n Fx$.

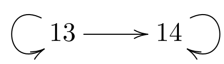

Remember that a number is determinately* small at a point iff you can't get to a smaller number by following the arrows. So, for example, the left node sees a node (itself) in which 14 is not determinately* small, and can see a node in which it is (the right node.) Thus, at the left node it is vague whether 14 is determinately* small, so 'determinately* small' has a borderline case. However the vagueness of 'determinately* small' is not determinate because the left node sees a node in which 'determinately* small' is completely precise: the right node.

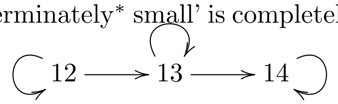

Here at the leftmost node it is vague whether 13 is determinately* small: it sees a world where it isn't (itself) and a world where it is (the middle node.) At the middle node it's vague whether 14 is determinately* small (see above). So at the leftmost node we have $\Delta(\nabla\Delta^* S(13) \vee \nabla\Delta^* S(14))$. Thus at every node the leftmost node sees 'determinately* small' has a borderline case, witnessed by 13 and 14 respectively. This gives a model for $\Delta\exists x\nabla\Delta^* Fx$ (and also $\Delta\exists x\nabla\Delta^n Fx$, for each n.)

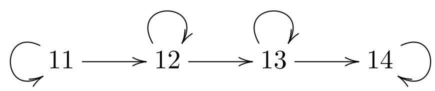

Just as before, it is vague at the leftmost node whether 12 is determinately* small. At every world it sees either 12 or 13 is a borderline case of determinate* smallness (see above), and at every node seen by a node seen by the leftmost either 12, 13 or 14 is a borderline case of determinate* smallness. So we have $\Delta\Delta(\nabla\Delta^* S(12) \vee \nabla\Delta^* S(13) \vee \nabla\Delta^* S(14))$. Thus this is a model for $\Delta\Delta\exists x\nabla\Delta^* Fx$ (and also $\Delta\Delta\exists x\nabla\Delta^n Fx$, for each n.) It should be clear how to carry on this series.

2.5 Nihilism*

One response to a Sorites paradox for a predicate of the form $\Delta^* Fx$ is simply to deny than *anything* satisfies $\Delta^* Fx$. I shall call this view 'nihilism*' since it mimics the nihilist response to the ordinary Sorites paradox.

It should be noted that the necessitation principle for Δ, a principle we have appealed to throughout this paper, already supplies counterexamples to nihilism* since, with a principle of infinitary conjunction introduction, we can easily show $\Delta^*(Fx \vee \neg Fx)$. Not only can we show that logical truths are determinate*, but we can also show things like $\Delta^*(\Delta Fx \rightarrow Fx)$, and indeed $\Delta^*(\Box Fx \rightarrow Fx)$ and $\Delta^*(KFx \rightarrow Fx)$ if we have necessity and knowledge operators in the language. Although certain logical and conceptual truths are among the determinate* propositions, some conceptual truths aren't. For example, if it is vague whether any foetus of a certain age is a person, then presumably it is is a conceptual truth one way or the other without being a determinate, and hence a determinate*, truth one way or the other.

Anyone who accepts necessitation for Δ has to admit a distinction between certain determinate* truths and others. I would be very skeptical that such a distinction would be a precise distinction.[17] And without a motivated precise distinction between the Δ^* truths and the rest, a response to the paradoxes of higher order vagueness is needed.

A more radical approach is to deny necessitation altogether. This is the strategy that Dorr seems to endorse in [8]. A standard way to model failures of necessitation is to introduce a distinction between normal and non-normal nodes in the kind of Kripke frames we have been considering. Such a move invites variants of the kinds of paradoxes we have been considering. For example, we could define an operator $\Delta_N p$ saying that p is true at all accessible *normal* worlds, and run the paradox for the operator $\Delta^*_N p$. Without appealing to this particular semantics for non-necessitatable operators, we must have some notion of normalcy which allows us to say things like '$\Delta p \rightarrow p$ is part of our logic and $\Delta p \rightarrow \Delta\Delta p$ isn't', i.e. something stronger than the mere truth or assertability of p, but weaker than the empty notion of determinate* truth.

[17] At least, there isn't any obvious precise criteria for distinguishing the two such as 'is a tautology' and so on.

3 Appendix

Here we demonstrate some relevant facts about the logic of higher order vagueness. Recall that:

Definition 3.0.2. *A v-frame is a triple* $\langle W, d(\cdot,\cdot), f(\cdot)\rangle$ *where* $\langle W, d\rangle$ *is a metric space, and* $f : W \to \mathbb{R}^+$ *obeys the following:*

(A) $\forall w, v \in W, |f(w) - f(v)| \leq d(w, v)$

A formula of propositional modal logic is valid on a v-frame $\langle W, d, f\rangle$ iff it is valid on the Kripke frame $\langle W, R\rangle$ where Rxy iff $d(x, y) \leq f(x)$.

Dorr [7] shows, translating into the terminology of v-frames, that B is not valid over the v-frame $\langle (1, 2), |x - y|, \frac{x}{3}\rangle$ although the weaker principles $p \to \Delta\neg\Delta\Delta\neg p$ and B^2 are valid in this frame. It is possible, however, to construct v-frames in which $p \to \Delta\neg\Delta^n\neg p$ is valid for no $n \in \mathbb{N}$. For example, let $W := \{0, 1\}$, $d(x, y) = |x - y|$, $f(0) = 1$ and $f(1) = \frac{1}{2}$.

What is the logic of v-frames? Clearly every v-frame generates a corresponding reflexive Kripke frame, so the logic of v-frames contains KT. One might have hoped that every reflexive Kripke frame could be generated from a v-frame this way ensuring a logic of exactly KT. This reduces to the question of whether every reflexive digraph can be embedded into a metric space in such a way that there is a a closed ball around each node that contains all and only those nodes it can see. Unfortunately this does not hold:

Fact: Suppose $\mathcal{F}$ is a Kripke frame based on a v-frame. If $\mathcal{F}$ contains a cycle, it contains a 2-cycle.

To see this suppose that $\langle a_0, \ldots, a_n\rangle$ is a cycle in $\mathcal{F} = \langle W, R\rangle$ where $n > 2$. For convenience let $a_i = a_j$ where $j = i \bmod (n + 1)$ for $i > n$. Now suppose that that $\neg Ra_{i+1}a_i$ for every i. Since for each i Ra_ia_{i+1} we know that $d(a_i, a_{i+1}) \leq f(a_i)$ in the corresponding v-frame. We also know that $f(a_i) < d(a_{i-1}, a_i)$ since $\neg Ra_ia_{i-1}$. Thus for each i, $d(a_i, a_{i+1}) \leq f(a_i) < d(a_{i-1}, a_i)$, so $f(a_n) < d(a_{n-1}, a_n) \leq f(a_{n-1}) < \ldots \leq f(a_{-1}) = f(a_n)$, i.e. $f(a_n) < f(a_n)$ which is a contradiction. So for some i, Ra_ia_{i+1} and $Ra_{i+1}Ra_i$.

v-frames thus have more structure than reflexive frames. However, it turns out this does not make a difference to the logic:

Theorem 3.1. Completeness. *A set* Σ *is valid on every v-frame iff it's members are theorem's of* KT.

Proof. Suppose that Σ is a KT-consistent set of formulae. Then Σ is satisfiable on the canonical frame $\mathcal{F}$. $\mathcal{F}$ may contain cycles without 2-cycles, so we cannot yet infer that Σ is satisfiable on some v-frame. However we may construct a frame from $\mathcal{F}$, with all the cycles ironed out, that is equivalent to a v-frame.

Let a_0 be a maximal KT-consistent set containing Σ. We may assume that a_0 is a root of $\mathcal{F}$ (if it isn't take the generated subframe around a_0 and work with that instead.) Define $\mathcal{F}^+ := \langle W^+, R^+\rangle$ as follows

- $W^+ := \{s \mid s \text{ a path in } \mathcal{F} \text{ such that } s_0 = a_0\}$

- $R^+ := \{\langle s,t\rangle \mid |t| = |s|+1 \text{ and } s_i = t_i \text{ for } i \leq |s| \text{ or } s = t\}$

Claim: $f(a_0,\ldots,a_n) = a_n$ is a bounded morphism from $\mathcal{F}^+$ to $\mathcal{F}$.

(1) Suppose R^+st. If $s = t$ then $f(s) = f(t)$ so $Rf(s)f(t)$ since R is reflexive. If $|t| = |s|+1$ then $Rf(s)f(t)$ since t and s are paths.

(2) Suppose $Rf(s)f(t)$. We want to find a u such that R^+su and $f(u) = f(t)$. If $f(t) = f(s)$ let $u = s$. Otherwise, let $u = \langle s, f(t)\rangle$.

Since anything valid on $\mathcal{F}^+$ is valid on every bounded morphic image of $\mathcal{F}^+$ (see for example [2]) it follows that Σ is satisfiable on $\mathcal{F}^+$. Now we construct our v-frame as follows:

- We begin by defining distance between adjacent points. If R^+st then $e(s,t) = e(t,s) = \frac{1}{2^{|s|}}$. Always fix $e(s,s) = 0$
- $d(s,t) := inf\{\sum_{i=0}^{n} e(p_i,p_{i+1}) \mid p_0 = s, p_n = t, p$ a path in the symmetric closure of $\mathcal{F}^+\}$
- $f(s) := \frac{1}{2^{|s|}}$

It is now easy to check that $\langle W^+, d, f\rangle$ is a v-frame and that R^+st iff $d(s,t) \leq f(s)$. $\square$

Cian Dorr has pointed out to me that the constraint (A) on v-frames does not play much of a role in the proof of Theorem 4.1. This allows us to prove a slightly more general result:

Definition 3.1.1. *a* **difference measure** *is a function $g : \mathbb{R}^2 \to \mathbb{R}$ such that:*

- *g is continuous in both arguments.*
- *$g(x,x) = 0$*
- *$g(x,y) = g(y,x)$ (this constraint is optional in what follows.)*

For a given difference measure, g, a g-frame is a triple $\langle W, d(\cdot,\cdot), f(\cdot)\rangle$ where $\langle W, d\rangle$ is a metric space, and $f : W \to \mathbb{R}$ such that:

(A′) $\forall w,v \in W, g(f(w), f(v)) \leq d(w,v)$

Corollary 3.2. *For any difference measure g, the logic of g-frames is* KT.

Proof. Note that for any positive a there is a $b < a$ such that $g(a,b) \leq a$ since $g(a,a) = 0$ and g is continuous in both arguments. For any a pick a unique such b, a_g (choice.)

Now modify the construction in Theorem 4.1 as follows.

- Fix $e(s,s) = 0$ for every s.
- Let $e(\langle a_0\rangle), t) = e(t, \langle a_0\rangle) := 1$ for $t \neq \langle a_0\rangle$ such that $R^+\langle a_0\rangle, t$.
- Suppose that $e(s,t) = e(t,s) = a$ has already been defined for R^+st, and suppose that R^+tu $t \neq u$. Define $e(t,u) = e(u,t) = a_g$.

- $d(s,t) := inf\{\sum_{i=0}^{n} e(p_i, p_{i+1}) \mid p_0 = s, p_n = t, p$ a path in the symmetric closure of $\mathcal{F}^+\}$
- $f(s) := sup\{e(s,t) \mid R^+st\}$.

□

The interest in this generalization is that one might think that it becomes much harder for two points to differ on the interpretation of 'determinately' the closer together they are. Perhaps it is not the difference between $f(w)$ and $f(v)$ that must be less than $d(w,v)$ but the difference between their ratios, or some other such g.

3.1 Further restrictions

The proof of completeness and the counterexamples to the B^n principles relied heavily on our considering slightly artificial frames that were based on metric spaces that either aren't dense, or have points with zero accessibility range. A natural class of v-frames to consider are those based on metric spaces of the form $\mathbb{R}^n$ where $f(a) > 0$ for all $a \in \mathbb{R}^n$. In these frames whenever x can see y, there is a path back from y to x, even though there are frames invalidating B^n for each $n \in \omega$ (i.e. there is no upperbound on how long these paths might be.) This is worrying since this means that $\Delta\Delta^*p \vee \Delta\neg\Delta^*p$ is valid over these frames.

We can express something like this principle in modal logic. I'll call it B^*.

$$\mathsf{B}^*\text{: } \Delta(p \to \Delta p) \to (\neg p \to \Delta\neg p) \tag{8}$$

B^* is valid in the class of v-frames just described. In the presence of KT, B^* defines what I shall call 'the backtrack principle'.

> Whenever Rxy there exists $z_1, \dots, z_n$ such that (a) $z_1 = y$, $z_n = x$ and Rz_iz_{i+1} for $1 \leq i < n$ and (b) Rxz_i for $1 \leq i \leq n$. (9)

Proof. We shall show that $(\Delta(p \to \Delta p) \wedge \neg\Delta\neg p) \to p$ defines the requisite property. Suppose the reflexive frame $\mathcal{F} = \langle W, R\rangle$ has the backtrack property. Now suppose $x \Vdash (\Delta(p \to \Delta p) \wedge \neg\Delta\neg p)$. The second conjunct ensures that there is a y such that Rxy and $y \Vdash p$. Since $\mathcal{F}$ has the backtrack property there is a finite path back from y to x, $z_1, \dots, z_n$, which x can see. Since $x \Vdash \Delta(p \to \Delta p)$ each $z_i \Vdash p \to \Delta p$. Since $z_1 = y$ and $y \Vdash p$, $y \Vdash \Delta p$ - by induction we can see that $z_i \Vdash p$ for each i which means $z_n = x \Vdash p$ as required.

For the other direction suppose, for contradiction, that $\mathcal{F} \models (\Delta(p \to \Delta p) \wedge \neg\Delta\neg p) \to p$ but $\mathcal{F}$ lacks the backtrack property. This means that for some x and y, Rxy but there is no path back from y to x which x can see. Define the following valuation on $\mathcal{F}$: $w \Vdash p$ iff there are $z_1, \dots z_n$ such that (1) $z_1 = y$, Rz_nw and Rz_iz_{i+1} for $1 \leq i < n$ and (2) Rxz_i for $1 \leq i \leq n$. Certainly if x had this property then $z_1, \dots, z_n, x$ would be a path back to x which x can see, so $x \nVdash p$. However $x \Vdash \Delta(p \to \Delta p)$ since if Rxw and $w \Vdash p$ then there is a path from y to w satisfying (1) and (2): $z_1, \dots z_n$. Furthermore, for any world that

w sees, w', $z_1, \dots, z_n, w$ will be a path from y to w' satisfying (1) and (2), since Rxw.

□

Is KTB^* the modal logic of these v-frames? We start with a negative result: KTB^* is not sound and *strongly* complete with respect to *any* class of frames. I.e. there is no class of frames, $\mathcal{C}$, such that a set is KTB^* consistent iff it's satisfiable on a frame in $\mathcal{C}$. The following also shows it is neither canonical nor compact.

Proof. To show this we shall show there is a KTB^*-consistent set of sentences which is unsatisfiable on every frame validating KTB^*.

Let $\Sigma := \{p, \neg\Delta\neg q\} \cup \{\Delta(q \to \Delta^n \neg p) \mid n \in \omega\}$. If Σ were KTB^*-inconsistent some finite subset would be KTB^*-inconsistent (since proofs are finite.) We shall show that for every $m \in \omega$, $\Sigma_m := \{p, \neg\Delta\neg q\} \cup \{\Delta(q \to \Delta^n \neg p) \mid n \in m\}$ is KTB^*-consistent. Σ_m has a KTB^*-model: $\langle m+1, R\rangle$ where Rxy iff $x = 0$ or $x > 0$ and $|x - y| \leq 1$. 0 can see m and there is a finite m length path back from m to 0 that 0 can see but no shorter path. Let q be true only at m and p only at 0.

However, if $\mathcal{F}$ validates KTB^* then $\mathcal{F}$ has the backtrack property so at no point of $\mathcal{F}$ is every member of Σ true: if $x \Vdash \neg\Delta\neg q$ then x sees some $y \Vdash q$. By the backtrack property there is a path $z_1, \dots, z_n$ back to x which x can see, so $\Delta(q \to \Delta^{n+1}\neg p)$ cannot be true at x if $x \Vdash p$.

□

However there is a positive result, namely that KTB^* is sound and complete over the class of reflexive frames with the backtrack property. For this result I refer the reader to [1], who shows that KTB^* has the finite model property.

Theorem 3.3. *If ϕ is KTB^*-consistent then it is satisfiable on a finite reflexive frame with the backtrack property.*

The question whether KTB^* the logic of v-frames over $\mathbb{R}^n$ in which $f(x) > 0$ remains open.

3.2 B entails that a conjunction of determinate truths is determinate

The aim is to show distributivity within the modal logic KB with infinitary conjunction (C1-C3 below.) For convenience I shall introduce an operator $\Diamond p := \neg\Delta\neg p$

D. $\bigwedge_{i<\omega} \Delta\phi_i \to \Delta \bigwedge_{i<\omega} \phi_i$.

KB

K $\Delta(\phi \to \psi) \to (\Delta\phi \to \Delta\psi)$

B $\phi \to \Delta\Diamond\phi$

Nec. if $\vdash \phi$ then $\vdash \Delta\phi$

C1. $\bigwedge_{i<\omega} \phi_i \to \phi_n$ for each $n < \omega$.

C2. $\bigwedge_{i<\omega}(\phi_i \to \psi_i) \to (\bigwedge_{i<\omega} \phi_i \to \bigwedge_{i<\omega} \psi_i)$.

C3. If $\vdash \phi_i$ for each $i < \omega$, $\vdash \bigwedge_{i<\omega} \phi$.

Claim: D is independent of K (and KT) + C1-C3.

Construct a Montague-Scott frame as follows (see [5]): Let $\mathcal{W} := \mathbb{N}$ and for each world $w \in \mathcal{W}$ let the necessary propositions at w, $N(w)$, be the cofinite subsets of $\mathcal{W}$ (if we are trying to model T as well we let $N(w) := \{X \cup \{w\} \mid X \text{ is cofinite}\}$.) Then $\langle \mathcal{W}, N\rangle$ satisfies:

1. $\mathcal{W} \in N(w)$ for all $w \in \mathcal{W}$
2. If $X, Y \in N(w)$ then $X \cap Y \in N(w)$
3. (For T) If $X \in N(w)$ then $w \in X$.

Thus our frame models K (/KT) including C1-C3. However it does not model D, as can be seen by letting $[\![p_i]\!] := \{n \in \mathbb{N} \mid n > i\}$ (for KT: $\{n \in \mathbb{N} \mid n > i\} \cup \{0\}$, allowing D to fail at 0.) On the other hand any Kripke frame (reflexive Kripke frame) will validate K (KT) along with D.

Although the distributivity of infinite conjunctions over Δ is independent of K, the distributivity of $\Diamond$ over infinite conjunctions, perhaps surprisingly, is not independent in this way and can be show given just some relatively uncontroversial principles governing infinite conjunction.

Lemma 3.4. $\Diamond \bigwedge_{i<\omega} p_i \to \bigwedge_{i<\omega} \Diamond p_i$

Proof. First note that $\Diamond \bigwedge_{i<\omega} p_i \to \Diamond p_j$ for each j, by C1 and the background modal logic of K. Then by C3 and then C2 we can infer $\Diamond \bigwedge_{i<\omega} p_i \to \bigwedge_{i<\omega} \Diamond p_i$. □

Theorem 3.5. *Although D is independent of K (and KT) it is not independent of, and is in fact entailed by, KB (and thus KTB.)*

Proof. B directly gives us:

$$\bigwedge_{i<\omega} \Delta p_i \to \Delta\Diamond \bigwedge_{i<\omega} \Delta p_i \tag{10}$$

We may also infer from our lemma that

$$\Delta\Diamond \bigwedge_{i<\omega} \Delta p_i \to \Delta \bigwedge_{i<\omega} \Diamond\Delta p_i \tag{11}$$

by applying necessitation and the K principle. Finally we have

$$\Delta \bigwedge_{i<\omega} \Diamond\Delta p_i \to \Delta \bigwedge_{i<\omega} p_i \tag{12}$$

because we have $\Diamond\Delta p_i \to p_i$ for each i by B. So by C3 and then C2 we get $\bigwedge_{i<\omega} \Diamond\Delta p_i \to \bigwedge_{i<\omega} p_i$, and by necessitation and K that gives (12).

But (10), (11) and (12) give distributivity. □

It should be noted that this argument does not appeal to any characteristically classical principles. Indeed this argument can be carried out provided one has the following rules of inference as primitive or derived:

→1. $\phi, \phi \to \psi \vdash \psi$

→2. $\phi \to \psi, \psi \to \chi \vdash \phi \to \chi$

→3. $\phi \to \psi \vdash \neg\psi \to \neg\phi$

and where KB is understood to contain K and the axioms $\phi \to \Delta\Diamond\phi$ and $\Diamond\Delta\phi \to \phi$.[18]

References

[1] R.A. Benton. A Simple Incomplete Extension of T which is the Union of Two Complete Modal Logics with F.M.P. *Journal of Philosophical Logic*, 31(6):527–541, 2002.

[2] P. Blackburn, M. De Rijke, and Y. Venema. *Modal logic*. Cambridge University Press, 2002.

[3] S. Bobzien. Higher-order vagueness, radical unclarity, and absolute agnosticism. *Philosopher's Imprint*, 10(10).

[4] J.A. Burgess. The sorites paradox and higher-order vagueness. *Synthese*, 85(3):417–474, 1990.

[5] B.F. Chellas. *Modal logic: an introduction*. Cambridge Univ Pr, 1980.

[6] C. Dorr. Vagueness without ignorance. *Philosophical Perspectives*, 17(1):83–113, 2003.

[7] C. Dorr. How Vagueness Could Cut Out At Any Order. *Unpublished*, 2010.

[8] C. Dorr. Iterating Definiteness. *Cuts and Clouds*, 1(9):550–586, 2010.

[9] M. Eklund. What vagueness consists in. *Philosophical Studies*, 125(1):27–60, 2005.

[10] D.G. Fara. An anti-epistemicist consequence of margin for error semantics for knowledge. *Philosophy and Phenomenological Research*, 64(1):127–142, 2002.

[18]Without the second version of B the most we could prove without double negation elimination would be things of the form $\Diamond\Delta\neg\phi \to \neg\phi$.

[11] D.G. Fara. Gap principles, penumbral consequence, and infinitely higher-order vagueness. *Liars and Heaps: New Essays on Paradox*, pages 195–222, 2003.

[12] H.H. Field. *Saving truth from paradox*. Oxford University Press, USA, 2008.

[13] K. Fine. The Impossibility of Vagueness. *Philosophical Perspectives*, 22(1):111–136, 2008.

[14] A. Mahtani. Can Vagueness Cut Out at Any Order? *Australasian Journal of Philosophy*, 86(3):499–508, 2008.

[15] D. Raffman. Demoting higher-order vagueness. *Cuts and Clouds*, pages 509–522, 2010.

[16] R.M. Sainsbury. Concepts without boundaries. *Vagueness: a reader*, pages 251–264, 1996.

[17] T. Williamson. *Vagueness*. Routledge, 1994.

[18] T. Williamson. On the structure of higher-order vagueness. *Mind*, 108(429):127, 1999.

[19] C. Wright. Is higher order vagueness coherent? *Analysis*, 52(3):129, 1992.

[20] C. Wright. The illusion of higher-order vagueness. *Cuts and Clouds*, pages 523–549, 2010.

[21] E. Zardini. Squeezing and stretching: How vagueness can outrun borderlineness. In *Proceedings of the Aristotelian Society (Hardback)*, volume 106, pages 421–428. John Wiley & Sons, 2006.