Luminosity and Vagueness*

Elia Zardini

Northern Institute of Philosophy

Department of Philosophy

University of Aberdeen

elia.zardini@abdn.ac.uk

July 17, 2012

1 Stage-Setting

Say that one is *in a position to know* that P iff [P and, in order for one to (come to) know that P, one only needs to (come to) believe[1] that P on grounds one has *already*

*Earlier versions of the material in this paper have been presented in 2008 at the Arché Audit and at the Arché Basic Knowledge Seminar (University of St Andrews); in 2009, at the Formal Epistemology Project Research Seminar (University of Leuven), at a research seminar at the University of Aarhus, at the 5th SOPHA Congress (University of Geneva), at the COGITO Epistemology Seminar (University of Bologna) and at the NIP Formal Epistemology Seminar (University of Aberdeen); in 2010, at the 2nd Workshop on *Vagueness and Physics, Metaphysics, and Metametaphysics* (University of Barcelona). I'd like to thank all these audiences for very stimulating comments and discussions. Special thanks go to Derek Ball, Stew Cohen, Annalisa Coliva, Laura Delgado, Dylan Dodd, Julien Dutant, Paul Égré, Gabriel Gomez, Chris Kelp, Hannes Leitgeb, Dan López de Sa, Sebastiano Moruzzi, Sven Rosenkranz, Ian Rumfitt, Johanna Seibt, Jason Stanley, Brian Weatherson, Robbie Williams, Crispin Wright and an anonymous referee. I'm also very grateful to the editors Philip Ebert and Martin Smith, who made extremely useful suggestions for improving both the content and the presentation of the paper (as well as for reducing the rather improbable length it had assumed throughout the years). But my greatest intellectual debt here is obviously to Tim Williamson, whose luminous writings on the topic have provided me with a constant source of inspiration. Although I work my way through to a view somewhat opposite to his, I cannot overstate how much this paper owes to the sharpness of his set-up and the lucidity of his arguments. In writing the paper, I have benefitted, at different stages, from an AHRC Postdoctoral Research Fellowship and a UNAM Postdoctoral Research Fellowship, as well as from partial funds from the project FFI2008-06153 of the Spanish Ministry of Science and Innovation on *Vagueness and Physics, Metaphysics, and Metametaphysics*, from the project CONSOLIDER-INGENIO 2010 CSD2009-00056 of the Spanish Ministry of Science and Innovation on *Philosophy of Perspectival Thoughts and Facts* (PERSP) and from the European Commission's 7th Framework Programme FP7/2007-2013 under grant FP7-238128 for the European Philosophy Network on *Perspectival Thoughts and Facts* (PETAF).

[1]The issue of what confidence requirements (if any) there are on knowledge will loom large in this paper (see especially section 9). However, when the specifics of such requirements are immaterial to the

available].[2] Now, it is a traditional philosophical idea that certain domains of facts are *fully open to our view*, and it is quite natural to gloss that idea by precisely saying that *one is always in a position to know* that the facts—whatever they are—belonging to those domains obtain. This paper is devoted to defending that idea, in particular against the threat represented by certain considerations concerning *reliability*, and to so doing by exploiting the *vagueness* pervading those domains: perhaps surprisingly, the paper argues that careful heeding of vagueness and its phenomena, far from forcing new and surprising limits on our knowledge, actually removes one of the main barriers—unreliability—often thought to stand in its way.

The shift from knowing to being in a position to know is already enough to dispose of many cheap objections to idea that certain domains of facts are fully open to our view. However, Conee [2005] raises a subtler problem for it hinging on the correct remark that it is possible that, for example, one feels cold while one does not know that one feels cold because, having received misleading evidence to the effect that one after all does not feel cold, one's justification for believing that one feels cold is defeated. In such a case, it is arguable that, even if one came to believe that one feels cold on the grounds that are already available to one, one would still not know that one feels cold (one would still need to acquire further grounds enabling one to discount the misleading evidence)—and so it is arguable that, after all, one is not in a position to know that one feels cold, with the point obviously generalising to just about any candidate for being a domain of facts fully open to our view.[3] This objection from defeat however simply exemplifies the *much more general* trouble that, for just about every fact, just about any subject might not be in an *adequate epistemic position* with respect to that fact. For example, it is possible that one feels cold while one does not know that one feels cold because, having taken a pill that makes one believe that one feels cold on the basis of thermal sensations of warm and makes one not believe that one feels cold 99 out of 100 times at which one feels cold, one's belief about one's feeling cold is not reliable. In such a case, it is arguable that, even if one came to believe that one feels cold on the grounds that are already available to one, one would still not know that one feels cold (one would still need to improve the reliability of one's belief about one's feeling cold that has been so dramatically lowered by the ingestion of the pill)—and so it is arguable that, after all, one is not in a position to know that one feels cold. If the idea that certain domains of facts are fully open to our view is to have

point at hand (as happens at this juncture in the text), I'll for simplicity acquiesce in the traditional view that what is required for knowledge at the level of confidence is belief. If one is persuaded by the kind of example put forth in Radford [1966], pp. 2–3 (or by others), one should strictly speaking replace 'one believes that P' and its relatives with whatever is one's favoured candidate for being what is required for knowledge at the psychological, non-epistemic level.

[2] Throughout, I use square brackets to disambiguate constituent structure in English.

[3] It is actually not completely clear that it is equally arguable that, after all, the fact that one feels cold is not fully open to one's view. If it is not equally arguable, that would show that the notion of being fully open to one's view involved in the traditional philosophical idea under consideration is more loosely connected than is usually assumed with epistemic notions such as that of being in a position to be known. However, merely to keep things simple, throughout I'll assume that there is no gap between a fact's being fully open to *our view* and its being fully open to *our knowledge* (where the latter, as I'm about to argue in the text, does not require that the fact be fully open to just about *everyone's* knowledge).

any chance, clearly all such possibilities of an epistemically decayed subject have to be screened off and attention has to be restricted to what we'll call '*optimal* subjects' (vague as that notion may be). In the following, we'll assume this to be so.

Assuming the familiar apparatus of possible-world semantics, we can define a notion, *luminosity*, which will help to articulate further this traditional philosophical idea, while also providing a simple and neat framework in which the argument of concern for this paper can be sharply presented and investigated. We proceed stepwise by introducing the following notions:

Definition 1. A *case* is a centred world—that is, a possible world with a designated subject (referred to by the pronoun 'one') and time (referred to by the present tense).

Definition 2. A *condition* is a subject- and time-sensitive proposition.

Definition 3. A condition $\langle C \rangle$[4] is *luminous* iff, for every case α, if $\langle C \rangle$ is true at α, then in α one is in a position to know that C.

A set of luminous conditions thus seems to correspond reasonably well to a domain of facts fully open to our view. The rest of this paper will be concerned with a very interesting argument (henceforth, 'the *AL-argument*'), first put forth by Tim Williamson (see Williamson [1995]—where however only an importantly different variant is developed—and then Williamson [1996]; Williamson [2000], pp. 93–113), to the effect that almost no interesting condition is luminous in this sense.

The argument takes as paradigmatic target the condition ⟨One feels cold⟩ and proceeds by considering the series of cases that correspond to the possible situation characterised by the following three features:

(POSNEG) One feels very cold at dawn and gradually warms up until one feels very warm at noon;

(FORM) Throughout the process, one is steadily attending to one's own thermal sensations and steadily considering, on this basis, whether one feels cold, forming, with regard to this matter and on this basis, whatever confidence one needs to form in order to achieve knowledge;

(CONF) While at dawn one is extremely confident that one feels cold and not confident at all that one does not feel cold, one gradually loses confidence that one feels cold and gradually acquires confidence that one does not feel cold, until at noon one is extremely confident that one does not feel cold and not confident at all that one feels cold.

Let t_1, t_2, $t_3 \dots t_{1,000,000,000}$ be distinct equidistant times from dawn to noon, and, for every n,[5] let α_n be the case one is in at time t_n.

[4] Throughout, I use '$\langle \varphi \rangle$' as a singular term referring to the condition expressed by φ.

[5] Sometimes (as here) I will implicitly assume that the range of 'n' and of its like is restricted to the set $\{x : 1 \leq x \leq 1,000,000,000\}$ if that's clear enough from the context.

Before turning to the AL-argument itself, a couple of comments on (FORM) and (CONF) are in order. Firstly, it will be a central issue in this paper to specify the notion of *confidence* employed in (FORM) and (CONF) (I'll try do so in section 5 by introducing the notion of doxa and contrasting it with the more usual notion of credence, and I'll argue in section 9 that reflection on the nature and properties of doxa goes a long way towards undermining the AL-argument). In the meanwhile, we'll rely on an intuitive understanding of the notion, according to which one's confidence that C is the degree of strength to which one believes that C.

Secondly, as we'll see, the AL-argument will attack luminosity by claiming that, in our situation, for some n, in α_n one feels cold but in α_n one does not *know* that one feels cold. Of course, that claim engages with luminosity only in the presence of (FORM), which in effect allows us to obliterate the distinction between knowing and being in a position to know. We've already observed that, in order for the notion of luminosity to capture the idea that certain domains of facts are fully open to our view, the domain of cases in definition 3 should be restricted to those whose subject is optimal. Thus, also the subject of the situation described by the AL-argument should be assumed to be optimal. However, given (FORM), we now see that additional assumptions have to be made about such a subject. More specifically, (FORM) requires that one do what one psychologically needs to do in order to know whether one feels cold, and so rules out, for example, that, throughout the process, for lack of attention or whatever other reasons, one places no confidence at all in the condition ⟨One feels cold⟩.

This seldom-noticed strength of (FORM) raises the following question: if (FORM) rules out that excessively cautious pattern of confidence, why does it not equally rule out, as being *also* excessively cautious, the kind of pattern of confidence described by (CONF)? My own answer is that, contrary to what the question seems to assume, given the facts described by (POSNEG) the kind of pattern of confidence described by (CONF) has actually some claim to be the *best* kind of pattern of confidence that an optimal subject could exemplify: it's certainly the one that best proportions changes in the subject's confidence to changes in her thermal sensations, and that might suffice to give it an edge over all other kinds of patterns of confidence. Against this, one might worry about the fact that the patterns of confidence of the kind described by (CONF) force that one's confidence in $\alpha_{500,000,000}$ that one feels cold is, at best, ever so slightly higher than fifty-fifty. Assuming that in $\alpha_{500,000,000}$ one still feels cold and that knowledge requires confidence considerably higher than fifty-fifty, (CONF) would then violate (FORM): in $\alpha_{500,000,000}$ one would not be doing what one psychologically needs to do in order to know whether one feels cold—in particular, one would not be forming a high enough confidence. In other words, in $\alpha_{500,000,000}$ one would fail to know that one feels cold already *for purely psychological reasons*—exactly what (FORM) was supposed to rule out in the attempt at engaging with luminosity. In the following, I'll usually grant the appealing idea that (CONF) can be taken correctly to describe the kind of pattern of confidence that the optimal subject has even when (FORM) is enforced, but, as we've just seen, one should be careful not to make assumptions that are inconsistent with this concession (see section 9).

2 The Anti-Luminosity Argument

The AL-argument starts with what I think is its crucial step, the derivation of a premise consisting in the *knowledge margin-for-error* principle:

(KMAR) For every n, if in α_n one knows that one feels cold, then in α_{n+1} one feels cold

(see Williamson [1992] for an interesting study of some philosophical motivations and consequences of margin-for-error principles). The derivation of this premise runs as follows. Assume that in α_n one knows that one feels cold. Then in α_n one is reasonably *confident* that one feels cold (otherwise, one would not know that one feels cold).[6] Hence, by (CONF), in α_{n+1} one is at worst ever so slightly less confident that one feels cold. Now, does one feel cold in α_{n+1}? We argue by (classical) *reductio* that one does. For assume for *reductio* that in α_{n+1} one does not feel cold. Then one's still considerably high confidence in α_{n+1} that one feels cold is false and hence *mistaken*,[7] and this in turn implies that one's confidence in α_n is not *reliable*. This is so because a very similar confidence in a very similar condition formed on a very similar basis in a very similar case—i.e. the confidence one has in α_{n+1} in ⟨One feels cold⟩—is mistaken. But if in α_n one's confidence that one feels cold is not reliable, in α_n one arguably does not know that one feels cold, contrary to our original hypothesis. This is so because of the general principle that knowledge requires reliability, which we interpret in our framework as requiring that, for every n, in α_n one knows that one feels cold only if in α_n one's confidence that one feels cold is reliable.[8] Therefore, by (classical) *reductio*, we'd better say that in α_{n+1} one does feel cold. But α_n was completely arbitrary, and so we can conclude to (KMAR).

Now, if any condition stands a chance of being luminous, it would seem that ⟨One feels cold⟩ does. As applied to our situation, the luminosity of that condition entails:

(LUM) For every n, if in α_n one feels cold, then in α_n one is in a position to know that one feels cold.

But at this point trouble arises. For the net effect of (FORM) is that, for every n, in α_n one knows that one feels cold iff in α_n one is in a position to know that one feels cold. Hence, (LUM) and (FORM) together entail:

[6]The principle that knowledge requires *reasonably high confidence*, albeit possibly weaker than the principle that knowledge requires *belief*, might begin to appear dubious once one starts to think about the kind of example mentioned in fn 1 (I'll say more about the relationships between belief and confidence in sections 5 and 6). Although clearly relevant to our concerns, I shall ultimately remain neutral in this paper as to whether the principles that knowledge requires belief or reasonably high confidence hold.

[7]Throughout, the operative sense of 'mistake' and of its like is that of *objective* mistake, a failure located in the brute relation between a subject's confidence and reality rather than in the formation of the confidence in the light of the subject's epistemic position. There would be much more to say on this notion (see Zardini [2012a]), but for our purposes we can simply assume that a *sufficient* condition for a confidence that C to be mistaken is that it is not the case that C.

[8]In this paper, I will neither contest the general principle nor—save for the clarification to be made in section 6—the interpretation offered in the text of what the principle requires in our framework. I think both issues are worthy of further investigation.

(LUM$^+$) For every n, if in α_n one feels cold, then in α_n one knows that one feels cold.

This, together with (KMAR), yields the catastrophical *sorites* premise:

(SOR) For every n, if in α_n one feels cold, then in α_{n+1} one feels cold as well,

which in turn contradicts (POSNEG) by familiar reasoning (which is valid in most logics of some philosophical interest; I'll say a bit more about this in fn 34). Hence, (POSNEG), (FORM), (KMAR) and (LUM) are jointly inconsistent. Since (POSNEG) and (FORM) should be unproblematic, either (KMAR) or (LUM) has to give. But (KMAR) was supported by the appealing derivation given in the previous paragraph, derivation which relied only on (CONF) and on other seemingly plausible premises. There thus seems to be every reason to conclude against (LUM). This is the AL-argument.

3 Sorites

There is no denying the fact that, when first confronted with the AL-argument, one is apt to feel that it is illicitly trading on the *vagueness* of the target conditions—in our presentation, on the vagueness of the conditions that one feels cold, that one is in a position to know that one feels cold and that one knows that one feels cold (and possibly other conditions as well that also, if in a more implicit fashion, play a crucial role in the argument). Vagueness arguably presents itself in a *variety* of phenomena: in a companion to this paper (Zardini [2012a]), I have looked at the AL-argument with a focus on that phenomenon of vagueness consisting in the *presence of borderline cases* (see fn 17), while here I will focus instead on that phenomenon of vagueness consisting in the *apparent absence of sharp boundaries*. The latter phenomenon is also known as '*sorites susceptibility*' and consists in the fact that a vague expression *seems* to fail to draw a sharp boundary between its positive and negative cases. For example, it seems preposterous to think that there is a sharp boundary between the cases in which one feels cold and those in which one does not—that is, it seems preposterous to think that, for some n, in α_n one feels cold and in α_{n+1} one does not feel cold.

Associated with sorites susceptibility is the circumstance that vague expressions are quite generally known to be susceptible to *sorites paradoxes*, where a premise (the sorites premise) that seems to capture well the apparent absence of sharp boundaries of an expression would seem to have unacceptable consequences. For example:

(BSOR) For every n, if a man with n hairs is bald, so is a man with $n + 1$ hairs

seems to capture well the apparent absence of sharp boundaries in the case of 'bald'. Yet, as shown by a version of the *sorites paradox*, it would seem to clash with the fact a man with 1 hair is bald and a man with $1,000,000,000$ hairs is not bald.

We then have that sorites premises like (BSOR) plus a standard logic sufficiently strong to run the sorites paradox are a cocktail for disaster, and so either has to go. On most contemporary theories of vagueness, what goes is in effect the perfect truth of sorites premises, *prima facie* plausible[9] as they might have been. Save for the discussion in fn 34, we'll follow this contemporary wisdom here—in fact, we'll assume that the logic is fully classical. From such a perspective committed to the rejection of sorites premises, the question has then to arise as to why, given that its form is very similar to that of (BSOR), (KMAR) is not just another sorites premise, deriving whatever plausibility it enjoys from the same source as a standard sorites premise like (BSOR) does.

The worry can be made sharper. For, given the vagueness of 'feels cold', it is just to be expected that whatever plausibility is immediately lent by such vagueness to the relevant standard sorites premise will also be transmitted to (KMAR). To see this, reflect that, in the case of 'feels cold', the relevant standard sorites premise is (SOR), and that, given the factivity of knowledge, (KMAR) is a straightforward consequence of (SOR). Since, in virtue of the sorites susceptibility of the vague expression 'feels cold', (SOR) is immediately plausible, such plausibility should also be transmitted to a straightforward consequence of (SOR) such as (KMAR). Hence, (KMAR) derives *at least part* of its plausibility from the spurious plausibility of (SOR). Could it then not be that actually the plausibility of (KMAR) *wholly* relies on such shaky grounds?

4 Sharp Boundaries and Sharpenings

For one, Williamson shows awareness of this worry, and his reply to it is basically the following (see e.g. Williamson [2000], pp. 102–106). Call 'an (admissible) *sharpening*' any way of making precise vague expressions (say, 'bald' and 'hairy') which:

(i) *Respects clear positive and negative cases* (and so will verify the following statements: a man with 1 hair is bald; a man with 2 hairs is bald; a man with 1,000,000,000 hairs is not bald etc.);

(ii) *Respects* so-called '*penumbral connections*' (and so will verify the following statements: if a man with 250,000 hairs[10] is bald, so is a man with 249,999 hairs; it is not the case that a man with 250,000 hairs is both bald and not bald; if a man with 250,000 hairs is bald, a man with 250,000 hairs is not hairy etc.);

(iii) *Decides borderline cases* one way or the other (and so might for example verify the following statements: a man with 249,999 hairs is bald; a man with 250,000 hairs is not bald; a man with 250,001 hairs is not bald etc.)

[9] In the following, I'll often use 'plausible' and its relatives for '*prima facie* plausible' and its relatives.

[10] To make the reader appreciate the full force of penumbral connections—connections that hold even in the penumbral area constituted by borderline cases—I'm assuming (plausibly or not) that it is borderline whether a man with 250,000 hairs is bald or not.

(for an extensive use of this notion in the construction of a theory of vagueness see Fine [1975]).

Now, typically, when a vague expression is made precise any plausibility the sorites premise might have had *vanishes*. For example, we can sharpen ‘bald’ so as to apply to all and only those men with at most 249,999 hairs. Under this (sharpened) understanding of ‘bald’, (BSOR) loses all its plausibility and indeed has a clearly false instance (namely, ‘If a man with $249,999$ hairs is bald, so is a man with $250,000$ hairs’). Williamson then claims that, on the contrary, (KMAR) retains its plausibility even when ‘feels cold’ and ‘knows’ are sharpened, which he treats as good evidence that (KMAR) does not receive all of its plausibility in a spurious way from the vagueness of ‘feels cold’ and ‘knows’ (after all, if sharpenings eliminate the plausibility of standard sorites premises, they will also eliminate any plausibility of any principle which wholly derives from sorites susceptibility).[11]

However, in giving this reply, Williamson only considers sharpenings that sharpen ‘feels cold’ without sharpening ‘knows’ and *vice versa*. For a friend of (LUM) who wishes to evaluate the source of the plausibility of (KMAR), most if not all such sharpenings are of dubious relevance, since they make (LUM) itself lose much of its plausibility. Firstly, take sharpenings that sharpen ‘feels cold’ and let for example ‘feels cold at least to degree $\mathfrak{f}$’ (with ‘$\mathfrak{f}$’ ranging over maximally specific degrees that can be assigned to sensations of cold) be a sharpening of ‘feels cold’.[12] Then, even granting (LUM) itself, the consequent sharpening of (LUM) as ‘For every n, if in α_n one feels cold at least to degree $\mathfrak{f}$, then in α_n one is in a position to know that one feels cold’[13] is problematic for any (admissible) value of ‘$\mathfrak{f}$’, since it cannot be definitely the case that the domain of precise facts concerning feeling cold at least to degree $\mathfrak{f}$—as opposed to the vague feeling cold—is fully open to our view by applying the concept of feeling cold.[14] Notice also that, for high enough values of ‘$\mathfrak{f}$’, the consequent sharpening of (KMAR) as ‘For every n, if in α_n one knows that one feels cold, then in α_{n+1} one feels cold at least to degree $\mathfrak{f}$’ is arguably false. For suppose without loss of generality that $\mathfrak{d}$ is the minimum degree that is not an admissible

[11]I should add that to follow Williamson and use sharpenings as a first-pass diagnostics for detecting sorites premises by no means commits one to some kind of “*precisificational*” theory of vagueness, which actually both Williamson and I would reject. Thanks to an anonymous referee for urging this clarification.

[12]Admittedly, given the complexity of our concepts of feeling cold and knowing, this proposal for sharpening may seem rather artificial and gross (even more so for the other kind of sharpening to follow in the text). Still, it helps to illustrate my point in a perspicuous and concise way. I trust that the same point will apply, *mutatis mutandis*, for more realistic and complex kinds of sharpening.

[13]For reasons which we don’t need to go into here, sharpenings are best understood as *not reaching inside the scope of verbs of propositional attitudes*. In any event, the relevant sharpening of (LUM) that does so reach (‘For every n, if in α_n one feels cold at least to degree $\mathfrak{f}$, then in α_n one is in a position to know that one feels cold at least to degree $\mathfrak{f}$’) would be utterly implausible for any value of ‘$\mathfrak{f}$’, since, given the limits of our power of discrimination with respect to our own maximally specific thermal sensations, there is no plausibility to the idea that the domain of precise facts concerning feeling cold at least to degree $\mathfrak{f}$—as opposed to the vague feeling cold—is fully open to our view.

[14]Reason: if that were definitely the case, since in some cases in which it is borderline whether one feels cold one does definitely feel cold at least to degree $\mathfrak{f}$, in all those cases one would definitely be in a position to know that one feels cold, and so, by the definite factivity of being in a position to know, one would definitely feel cold, which contradicts the fact that those are cases in which it is borderline whether one feels cold.

value of '$\mathfrak{f}$' (i.e. the minimum degree such that, if one feels cold at least to that degree, one definitely feels cold). Also without loss of generality, let n be such that in α_n one feels cold exactly to degree $\mathfrak{d}$. Then, very plausibly, in α_n one knows that one feels cold. Moreover, in α_{n+1} one feels cold to degree $\mathfrak{f} - \mathfrak{e}$. By the properties of the reals and the definition of $\mathfrak{f}$, there is a high enough degree $\mathfrak{b} > \mathfrak{f} - \mathfrak{e}$ that is an admissible value for '$\mathfrak{f}$'. In α_{n+1} one does not feel cold at least to degree $\mathfrak{b}$, and so (KMAR) is false under such sharpening. Hence, even restricting our attention to the simple-minded sharpenings considered so far, there actually are sharpenings that make (KMAR) false.[15]

Secondly, take sharpenings that sharpen 'knows' and let for example 'knows at least to degree $\mathfrak{k}$' (with '$\mathfrak{k}$' ranging over maximally specific degrees that can be assigned to states of knowledge) be a sharpening of 'knows'. Then, even granting (LUM) itself, the consequent sharpening of (LUM) as 'For every n, if in α_n one feels cold, then in α_n one is in a position to know at least to degree $\mathfrak{k}$ that one feels cold' is problematic for high enough values of '$\mathfrak{k}$', since it cannot be definitely the case that the domain of facts concerning feeling cold is fully open to our precise knowing at least to degree $\mathfrak{k}$—as opposed to our vague knowing.[16] Notice also that, for low enough if not for all values of '$\mathfrak{k}$', the consequent sharpening of (KMAR) is arguably not definitely true. For suppose that $\mathfrak{b}$ is some such value and that in α_n one knows at least to degree $\mathfrak{b}$ that one feels cold. Then, if the consequent sharpening of (KMAR) were definitely true, since in α_n one does definitely know at least to degree $\mathfrak{b}$ that one feels cold, in α_{n+1} one would definitely feel cold, and so, very plausibly, one would still know at least to degree $\mathfrak{b}$ that one feels cold, thus yielding a sorites premise. Hence, even restricting our attention to the simple-minded sharpenings considered so far, there actually are sharpenings that make (KMAR) not definitely true.

Now, a sharpening can be inadmissible even if, relative to any *single* expression, it satisfies all of (i)–(iii). For example, a sharpening of 'bald' as 'has at most 250,000 hairs' may satisfy all of (i)–(iii) relative to the single expression 'bald', and yet be such as to sharpen 'hairy' as 'has at least 250,000 hairs': even though such a sharpening respects all the penumbral connections *internal* to 'bald', it does not respect the *external* penumbral connection between 'bald' and 'hairy' 'If a man with 250,000 hairs is bald, a man with

[15] Thanks to Philip Ebert and Martin Smith for pressing me on this last issue.

[16] Reason: if that were definitely the case, since in some cases in which it is borderline whether one feels cold one is definitely not in a position to know at least to degree $\mathfrak{k}$ that one feels cold, in all those cases one would definitely not feel cold, which contradicts the fact that those are cases in which it is borderline whether one feels cold. In this connection, it should be noted that the idea that certain domains of facts are fully open to our view has often been associated with the idea that our access to such facts is in some sense "*infallible*", which can in turn reasonably be understood to mean that one is in a position to know with *certainty* the relevant conditions. Although I'm actually sympathetic to this latter idea for at least some domains of facts, all the arguments in the paper will be neutral with respect to it: the arguments will only try to defend the former idea that, in the relevant cases, one always has *knowledge*, but they will be silent on what *strength* such knowledge has. In particular, when I say in this fn that, in some cases in which it is borderline whether one feels cold, one is not in a position to know at least to degree $\mathfrak{k}$ that one feels cold, that is not supposed to imply that, in such cases, one is not in a position to know with certainty that one feels cold: the degrees that measure the strength of one's knowledge are not the degrees that give rise to the different sharpenings of 'knows' (and of 'knows with certainty'). Thanks to Philip Ebert and Martin Smith for pointing out to me the need to mention the issue of certainty.

250,000 hairs is not hairy', and hence is inadmissible. The fact that, on such sharpening, 'Some men are both bald and hairy' is true clearly does not even begin to assuage the (right) doubts concerning the truth of that sentence. Now, one of the lessons of the previous two paragraphs is that, just as admissible sharpenings of 'bald' are only those that *coordinate* the extension of 'bald' and the extension of 'hairy' so that they are disjoint, the only admissible sharpenings for a friend of (LUM) are those that coordinate the extension of 'feels cold' and the extension of 'knows' so that they are connected in such a way that, for every α, one falls under the extension in α of 'feels cold' iff one falls under the extension in α of 'knows that one feels cold' (the left-to-right direction being given by (LUM) itself, the right-to-left direction being given by the factivity of knowledge, another penumbral connection). In effect, in its taking into account in the antecedent of a conditional absolutely every case in which one feels cold (and not just simply every case in which one definitely feels cold),[17] (LUM) aspires to be a penumbral connection, which, on every sharpening, has to hold non-vacuously in the borderline area no less than in the non-borderline area. For a friend of (LUM) then (LUM) acts as a penumbral connection that makes inadmissible the non-coordinated simple-minded sharpenings discussed in the previous two paragraphs.

Indeed, *from this perspective*, tables can be turned, for, if we now focus on the sharpenings that, in the manner just described, coordinate the extension of 'feels cold' and the extension of 'knows', we find that, far from continuing to support (KMAR), they all strip it of any plausibility. Indeed, on all such sharpenings, (KMAR) is clearly false, since the last case in which one feels cold will also be a case in which one knows that one feels cold, thereby verifying (KMAR)'s antecedent but falsifying its consequent. From this perspective, the clear falsity of (KMAR) on all the admissible sharpenings strongly suggests that (KMAR) does receive all of its plausibility in a spurious way from the vagueness of 'feels cold' and 'knows'.

I have taken the pain to stress, and emphasise once again to avoid misunderstandings, that this train of thought is, as a whole, only compelling *from the perspective of a friend of* (LUM). As we've seen, a friend of (KMAR) should agree that the only admissible sharpenings are those that coordinate *somehow* the extension of 'feels cold' and the extension of 'knows', but will then proceed to specify the nature of this coordination

[17] In Zardini [2012a], I consider theories of vagueness ('NFMU-theories') according to which, if it is borderline whether C, there is *no fact of the matter* as to (equivalently, it is *indeterminate*) whether C, and, if there is no fact of the matter as to whether C, it is *unknowable* whether C. I argue that, on such theories, principles like (LUM) already have to be restricted for independent reasons that *have nothing to do with our cognitive powers.* I then examine the most natural and promising restriction:

(DLUM) For every n, if in α_n it is determinately the case that one feels cold, then in α_n one is in a position to know that one feels cold,

arguing that, appearances notwithstanding, it does not fare better than (LUM) with respect to the AL-argument. In this paper, I wish to maintain neutrality as to whether an NFMU-theory of vagueness is correct, and will for simplicity focus on (LUM) rather than (DLUM)—obviously, any defence of (LUM) will *ipso facto* be also a defence of (DLUM) (I'll indicate in fn 37 how an important part of my defence can be modified so as to suit an NFMU-theory).

in a way alternative to that just offered on behalf of a friend of (LUM). However, even if the previous considerations do not vindicate *from a neutral perspective* the suspicion that (KMAR) receives all of its plausibility in a spurious way from the vagueness of 'feels cold' and 'knows', they do show that it does so *under a certain choice of coordinated sharpenings* (the one that a friend of (LUM) would make), and they do show that some choice or other of coordinated sharpenings has to be made. And that is enough to show that, absent reasons to discount that choice and to go instead for another quite specific and unobvious choice more hospitable to (KMAR), the attempt of acquitting (KMAR) from that suspicion by appeal to sharpenings does not succeed.

However, matters can hardly be left at that, since it is not the case that the AL-argument simply appeals to the *immediate plausibility* of (KMAR). Quite the contrary, as we've seen at the beginning of section 2, a significant part of the argument consists exactly in the *derivation* of (KMAR) from more fundamental principles concerning knowledge and reliability. We therefore still need to substantiate the suggestion that (KMAR) receives all of its plausibility in a spurious way from vagueness by showing where precisely the derivation in question goes astray. Doing so will allow us to understand how one can have reliable knowledge even at the penumbral limit with falsity.

5 Credence and Doxa

Before turning to that task, some words should be spent on the notion of "confidence" employed by the AL-argument. Even though it too can be associated with a degree scale, that notion should be sharply distinguished from the notion of *credence*. I think it should rather be identified, roughly speaking, with the notion of *degree of outright belief*,[18] which for short I'll call '*doxa*' (plural 'doxai'). Let me explain. I take credences to concern the degree to which one *expects* something to be the case, to correlate, given certain idealisations, with *betting behaviour* and to be normatively governed by the *classical* probability calculus. In contrast, I take doxai to concern the degree to which one *outright believes* something to be the case, to correlate, given certain idealisations, with *assertion* and *use in practical syllogism* and to be normatively governed by some *non-classical* probability calculus (for example, the Dempster-Shafer theory, see Dempster [1967]; Shafer [1976]).[19] Unlike expectation, outright belief concerns one's *investment in*, or *commitment*

[18]Henceforth, in order better to contrast it with graded properties of partially believing, I'll use 'outright belief' and its relatives to denote the all-or-nothing property of believing that until now I've been denoting with 'belief' and its relatives.

[19]The references in the text should be taken as merely suggestive. A notable feature of Dempster-Shafer functions is their failure to satisfy the *finite-additivity* axiom $\mathfrak{P}(X \vee Y) = \mathfrak{P}(X) + \mathfrak{P}(Y)$ (with X and Y incompatible), and the examples to follow in the text (especially the kind of case (b)) will make clear that, insofar as they motivate a departure from classical probability measures, they motivate an adoption of non-additive (indeed, super-additive) measures. But it is very much open exactly which kinds of measures are best suited to doxai, and the Dempster-Shafer theory is only one (worthy) candidate for this role (and different candidates may be best for different purposes). I stress that all the discussion to follow could easily be adapted to other super-additive measures, such as e.g. the more constrained necessity functions of Dubois and Prade [1985].

to, *the truth* of a certain condition: hence, on the input side, the conditions which licence its being held are stronger than those for holding an expectation; correspondingly, on the output side, its being held licences stronger consequences than holding an expectation does (we'll see in some detail in this section what consequences all this has for reliability). Like expectation and many other things, outright belief can come in degrees, and so, just as it makes sense to talk of an apple as being red only to a certain degree, it also makes sense to talk of someone as outright believing a condition only to a certain degree.

Such a distinction between two different species of confidence is of course controversial and this is by no means the place to go into the details of its explanation and defence, but, to give a flavour of it, let me briefly mention a couple of kinds of cases in which doxai and credences can come dramatically apart:[20]

(a) *Doxai may be higher than the corresponding credences.* For example, given the usual quantum-mechanical considerations, the objective chance, in normal circumstances, of one's not being on Mars in the next second is only, say, .999999, and so should presumably be one's credence (provided that one knows about the objective chance but does not have inadmissible evidence about the future), while one's doxa that one will not be on Mars in the next second may well be maximum.

(b) Conversely, *credences may be higher than the corresponding doxai.* For example, if one does not have any specific insider information, one's credence that Obama has had eggs for breakfast this morning may be, say, .5 (ditto for one's credence that Obama has not had eggs for breakfast this morning), while one's doxa that Obama has had eggs for breakfast this morning may well be minimum (ditto for one's doxa that Obama has not had eggs for breakfast this morning)—one may well not believe at all that Obama has had eggs for breakfast this morning (ditto for one's not believing at all that Obama has not had eggs for breakfast this morning).[21]

Why has the "confidence" in the AL-argument to be interpreted as doxa rather than credence? Because, as Williamson [2000], pp. 98–99 himself illuminatingly points out, under the credence interpretation, the fact that a very similar confidence in a very similar condition formed on a very similar basis in a very similar case—i.e. the confidence one has in α_{n+1} in $\langle$One feels cold$\rangle$—is mistaken does not suffice to show that one's confidence in α_n is not reliable (in the sense of reliability in which a charge of unreliability puts the knowledgeability of one's outright belief into question). To take a particularly extreme

[20] In the following, I'll assume that the relevant epistemic subjects do indeed perfectly conform to whatever are the normative principles governing credences and doxai respectively.

[21] Note that while the kind of case (a) can easily be turned, by taking the negation of the condition in question, into a kind of case in which credences may be higher than the corresponding doxai (assuming, very plausibly, that credences obey the classical probability calculus and that doxai obey, among others, the principle that, if one's doxa in $\langle C \rangle$ is 1, one's doxa in the negation of $\langle C \rangle$ is 0), the kind of case (b) cannot so easily be turned, conversely, into a kind of cases in which doxai may be higher than the corresponding credences. This would require something along the lines of the very implausible principle that, if one's doxa in $\langle C \rangle$ is 0, one's doxa in the negation of $\langle C \rangle$ is 1 (indeed, (b) itself constitutes a counterexample to any principle along those lines).

example, consider a case λ in which one knows that all the 100 tickets of a fair lottery (with exactly one winner) have been bought by Italians. Suppose that, on this basis, in λ one has a 1 credence that the winner will be Italian, and that one indeed outright believes that the winner will be Italian. Suppose also that in λ it might easily have been the case that only 99 tickets of the lottery have been bought by Italians (the remaining one having been bought by a Frenchman), and that, in such a new case λ', although one has a .99 credence that the winner will be Italian on a very similar basis as that in λ, being (rightly or wrongly) cautious one no longer outright believes at all that the winner will be Italian. Suppose, finally, that in λ' the lottery is after all won by the Frenchman.

Very intuitively, in λ one does know that the winner will be Italian, despite the fact that a very similar credence in a very similar condition formed on a very similar basis[22] in a very similar case—i.e. the credence one has in λ' in ⟨The winner will be Italian⟩—is mistaken. That fact alone does not suffice to give rise to a charge of unreliability against the knowledgeability of one's outright belief in λ. Granted, it would if—contrary to what we supposed—in λ' one still outright believed, at least to a certain degree, that the winner will be Italian, but that just shows the necessity of appealing to doxai rather than credences.

Importantly, I don't think that the problem is essentially tied to kinds of cases in which one's credence is 1. For, broadly speaking, the problem arises because, while on

[22]The fact that in λ, but not in λ', one has *conclusive evidence* that the winner will be Italian might be thought to cast doubt on the claim that the credence one forms in λ is formed on a basis similar to the basis on which one forms the credence in λ'. Against this doubt, there are undoubtedly many respects in which there is indeed similarity in the two credences' bases: both are formed and calibrated in reaction to information about the (same) features of the lottery, both are formed and calibrated in reaction to information about the (similar) number of Italian participants, both are formed and calibrated according to standard principles of probabilistic reasoning etc. I take it as very plausible that these many respects of similarity do suffice for the kind of similarity at issue in questions of reliability and knowledge; if this is to be convincingly denied, specific reasons must be given as to why this is not the case. In this regard, I wish to note that the dissimilarity just mentioned—consisting in the fact that in λ, but not in λ', one has conclusive evidence that the winner will be Italian—cannot be taken as a good reason to deny that the credence one forms in λ is formed on a basis similar to the basis on which one forms the credence in λ', on pain of the resulting notion of reliability losing its stringency as a requirement on knowledge. Consider the simple variant of the example in which in λ it might easily have been the case that, again, only 99 tickets of the lottery have been bought by Italians (the remaining one having been bought by a Frenchman), but also that, in such a new case λ'', one has a 1 credence that the winner will be Italian, and that one indeed outright believes that the winner will be Italian. Suppose, finally, that in λ'' the lottery is after all won by the Frenchman. I take it as plausible that, in this variant of the example, in λ one does not know that the winner will be Italian, and does not know it because one's outright belief is not reliable. Although in λ one still has conclusive evidence that the winner will be Italian, the easy possibility of λ'' shows that one is not sensitive to the conclusive character of such evidence—one might easily react in the same way even if one did not have it, and indeed even if it were not the case that the winner will be Italian. But we could not use λ'' in this way to impugn the reliability of one's outright belief in λ if the fact that in λ, but not in λ'', one has conclusive evidence that the winner will be Italian were a good reason to deny that the credence one forms in λ is formed on a basis similar to the basis on which one forms the credence in λ''. Notice also that the feature of having conclusive evidence is completely absent in the next example to be introduced in the text. Thanks to Paul Égré for pressing me on these issues.

the one hand high credence does not guarantee high doxa,[23] on the other hand there may be a correlation between a certain threshold of credence and a substantially high doxa, indeed a knowledgeable one. In the kind of case just described, the threshold was 1, but these two features can be plausibly replicated also for lower thresholds. Here is one fairly general way of doing this. Coming back to a theme that has already emerged in the kind of case (a), consider a series of different quantum-like cases κ_0, κ_1, $\kappa_2 \ldots$, differing among themselves chiefly with respect to the threshold that they set for the (very high) objective chances, in normal circumstances, of one's not being on Mars in the next second (to fix ideas, let these chances be .999999, .9999999, .99999999 . . .). Suppose that circumstances are normal, that one knows that in effect the objective chance of one's not being on Mars in the next second is .999999 (.9999999, .99999999. . .), and that one indeed outright believes that one will not be on Mars in the next second. Suppose also that in κ_0 (κ_1, $\kappa_2 \ldots$) it might easily have been the case that the objective chance of one's not being on Mars in the next second is only .999998 (.9999998, .99999998. . .) and that, in such a new case κ'_0 (κ'_1, $\kappa'_2 \ldots$), although one has a .999998 (.9999998, .99999998. . .) credence that one will not be on Mars in the next second on a very similar basis as that in κ_0 (κ_1, $\kappa_2 \ldots$), being (rightly or wrongly) cautious one no longer outright believes at all that one will not be on Mars in the next second. The high similarity in the basis of one's credences can be ensured by supposing that in both κ_0 and κ'_0 (κ_1 and κ'_1, κ_2 and $\kappa'_2 \ldots$) one simply gets to know the chances from the testimony of a scientist one rightly and completely trusts. The dramatic divergence in doxastic investment between the twin cases κ_n and κ'_n can in turn be ensured by supposing that in both κ_0 and κ'_0 (κ_1 and κ'_1, κ_2 and $\kappa'_2 \ldots$) one knows that [the chance of one's not being on Mars is at least as high as .999999 (.9999999, .99999999. . .) iff one is in normal circumstances] (and that [one would invest a substantial amount of doxa in a condition to which one does not assign a 1 credence iff one knows that the credence's deviation from 1 is merely due to the usual quantum-like considerations applying in normal circumstances]).[24] Suppose, finally, that in κ'_0 (κ'_1, $\kappa'_2 \ldots$) one is after all teletransported to Mars in the next second.

Very intuitively, in κ_0 (κ_1, $\kappa_2 \ldots$) one does know that one will not be on Mars in the next second, despite the fact that a very similar credence in a very similar condition

[23] In a sense, this may generalise even to credences of degree 1. One way of seeing this is to think, beyond the kinds of cases (a) and (b), about certain kinds of cases involving infinity. For example, unless—somewhat implausibly—credences are taken to obey the *countable-additivity* axiom $\mathfrak{P}\left(\bigcup_{i\in\omega}(X_i)\right) = \sum_{i\in\omega}(\mathfrak{P}(X_i))$ (with X_i and X_j incompatible whenever $i \neq j$), one's credence that it is not the case that one's ticket will win in a fair lottery with denumerably many tickets may be 1 (since one's credence that one's ticket will win in a fair lottery with denumerably many tickets may be 0), but that need not apply to one's doxa (even if one's doxa that one's ticket will win in a fair lottery with denumerably many tickets may also be 0). This would require the very implausible principle mentioned at the end of fn 21.

[24] Why is it the case that in κ'_0 (κ'_1, $\kappa'_2 \ldots$) the chance of one's not being on Mars in the next second is only .999998 (.9999998, .99999998. . .)? Well, for example it might be the case that, as opposed to κ_0 (κ_1, $\kappa_2 \ldots$), in κ'_0 (κ'_1, $\kappa'_2 \ldots$) circumstances are not normal and one is used as a guinea pig for an indeterministic teletransporter that has an objective chance of .000002 (.0000002, .00000002. . .) of teletransporting one to Mars in the next second.

formed on a very similar basis in a very similar case—i.e. the credence one has in κ'_0 (κ'_1, $\kappa'_2 \ldots$) in ⟨One will not be on Mars in the next second⟩—is mistaken. That fact alone does not suffice to give rise to a charge of unreliability against the knowledgeability of one's outright belief in κ_0 (κ_1, $\kappa_2 \ldots$). Granted, it would if—contrary to what we supposed—in κ'_0 (κ'_1, $\kappa'_2 \ldots$) one still outright believed, at least to a certain degree, that one will not be on Mars in the next second, but that just shows the necessity of appealing to doxai rather than credences.

We can then conclude that the correct interpretation of the ambiguous notion of "confidence" in the AL-argument is in terms of doxai rather than credences. This conclusion highlights an important *commitment* of the argument: the psychological reality of a species of confidence different from credence, for, as we've seen, credences just do not support the reasoning about reliability that the argument requires. I myself have just argued in favour of the existence of a particular species of confidence different from credence, and so do not find that commitment in itself objectionable; from this perspective, the conclusion may appear to be overall a *friendly gift*, in that it gives the AL-argument the best chance to succeed. However, I actually mean to put forth the conclusion ultimately as a *Trojan horse*: as I'll explain at the beginning of section 9, doxai have formal properties that allow a straightforward modelling which conclusively establishes by everyone's lights the reliability of one's doxa that one feels cold throughout the situation described by the AL-argument.[25]

6 Reliable Knowledge at the Limit

It will prove useful, in order to appreciate the workings of my argument against (KMAR), to offer first a simpler version of the argument which focusses on outright beliefs, and then proceed to develop a more general and more powerful version which dispenses with outright beliefs and only works with doxai. Start by considering that, by the least-number principle, there is a last m such that in α_m one outright believes that one feels cold. Let b be such an m. Given (CONF), it is uncontroversial that $b < 1,000,000,000$. Hence, in α_{b+1} it is simply not the case that one outright believes that one feels cold.[26]

[25] In any event, for those that still think that the correct interpretation of the ambiguous notion of "confidence" in the AL-argument is in terms of credences rather than doxai, I hasten to add that, apart from the twist just foreshadowed in the text, all the arguments to follow cast in terms of doxai can easily be reformulated in terms of credences. Thanks to Sven Rosenkranz and an anonymous referee for urging this clarification.

[26] I emphasise once and for all that the existence of this and similar sharp boundaries follows simply from the least-number principle and similar logical or mathematical facts. As such, these boundaries do not imply any controversial "*threshold* view" to the effect that, for example, one believes that C iff one's doxa that C is above a certain (possibly only contextually determined) threshold. For example, the point just made in the text is consistent with the possibility that in α_{b-1} one does not outright believe that one feels cold, with the possibility that [one's doxa that one feels tired has the same degree as one's doxa in α_b that one feels cold but one does not outright believe that one feels tired], with the possibility that [in another situation one's doxa that one feels cold has the same degree as one's doxa in α_b that one feels cold but in that situation one does not outright believe that one feels cold], possibilities which are most

Return now to that part of the AL-argument which derives (KMAR) from more fundamental principles concerning knowledge and reliability. At the crucial junction, in the sub-argument for *reductio*, we are asked to assume that, for an arbitrary n, in α_n one knows (and so has a reasonably high doxa) that one feels cold, but that in α_{n+1} one does not feel cold. We are asked then to infer from the fact that, by (CONF), in α_{n+1} one has still a considerably high doxa—at worst *ever so slightly lower* than the one one has in α_n—that one's doxa in α_n is not reliable. The inference is supposed to be licenced by the consideration that a *very similar doxa* in a very similar condition formed on a very similar basis in a very similar case—i.e. the doxa one has in α_{n+1} in ⟨One feels cold⟩—is mistaken. But it should by now be clear that such a consideration concerning the high similarity in the doxai is at best of dubious relevance. For it is now open to the friend of (LUM) to retort that, for all the AL-argument has shown, n may well be b, in which case, although one's doxa in α_{n+1} is very similar to one's doxa in α_n, in α_{n+1}—as opposed to α_n—one does *no longer outright believe* that one feels cold. And if in α_{n+1} one does no longer outright believe that one feels cold, it is hard to see how the fact that in α_{n+1} one does not feel cold could give rise to a charge of unreliability against the knowledgeability of one's outright belief in α_n.

Strangely enough, Williamson for one seems to be aware of this problem but also seems to think that it can easily be dismissed by the following point:

> Even if one's confidence at t_n was just enough as to count as belief, while one's confidence at t_{n+1} falls just short of belief, what *constituted* that belief at t_n was largely misplaced confidence; the belief fell short of knowledge. (Williamson [2000], p. 97; the emphasis and the notational alterations are mine.)

This passage is very concise and there might be better interpretations of it than those I've been able to come up with, but, at least as I read it, I find it rather unconvincing. In its use of 'constitution'-talk, it seems to rely at least in part on the somewhat inchoate and speculative idea that the doxa one has in α_{n+1} can be identified with an equally sized portion of the doxa one has in α_n. Needless to say, that idea is by no means forced by the assumptions of the case: all one can say is that the doxa in α_{n+1} has a very similar degree to the doxa in α_n and is formed on a very similar basis, but no conclusion can be drawn about the identity of the former with portions of the latter.[27] Moreover, the idea itself is clearly a non-starter in many *inter-personal* cases (your degree of weight is very similar to mine, just ever so slightly lower: it certainly doesn't follow, unless we share substantial parts of our bodies, that your weight is identical with or constitutes a large portion of mine); and it is also clearly misguided in quite a few *intra-personal*, *inter-temporal* cases

---

or all inconsistent with the threshold view just sketched. Thanks to an anonymous referee for pressing me on this issue.

[27]This identity could of course be stipulated into the situation. However, such a stipulation implies that one's doxa is more open to the risk of being misplaced than it would be if the stipulated identity did not hold. Similarly to (albeit in a lesser measure than) the dramatic case discussed in section 1 (where one takes a pill that makes one vastly unreliable in one's belief that one feels cold), the stipulation thus implies that one is not an optimal subject, and so not relevant to issues of luminosity.

(my degree of income this year is very similar to my degree of income last year, just ever so slightly lower: it certainly doesn't follow, unless this year's salary was paid in advance as part of last year's salary in a double-counting way, that my income this year is identical with or constitutes a large portion of my income last year). A highly non-trivial argument is wanted to the effect that the required identifications in the case of doxai fare any better.[28]

These are however relatively minor skirmishes. For the main point remains that, if n is b, in α_n one outright believes that one feels cold while in α_{n+1} one does not outright believe that one feels cold: one's outright belief *perfectly tracks* the jump from feeling cold to not feeling cold and is thereby *never mistaken*. It is hard to see any plausible sense in which such an infallible outright belief could still be claimed to be unreliable.

One might rejoin that the compellingness of the point depends on tacitly assuming outright beliefs, rather than doxai, to be the *bearers of reliability and unreliability*. That assumption seems to me highly plausible and goes unchallenged by anything argued for in the AL-argument.[29] Be that as it may, I don't think that the move of assuming doxai to be the bearers of reliability and unreliability ultimately succeeds in defusing the style of argument I'm running. Firstly, I want to stress that it is very unclear that the crucial step in the AL-argument under discussion is any less fallacious if we assume doxai instead to be the bearers of reliability and unreliability. For consider that, exactly because doxai are somewhat functional to determine degrees and comparisons of one's doxastic investments, no matter how *similar* they can be *in degree* they can still be *related to two doxastically very different states*—outright belief and lack of outright belief—and so be themselves qualitatively very different in crucial respects. In particular, while it may well be accepted that a mistaken doxa as high as one's doxa in α_b would give rise to a charge of unreliability against the knowledgeability of a doxa had in similar enough cases, it may equally well be rejected that a mistaken doxa as high as one's doxa in α_{b+1} would also give rise to such a charge. After all, no matter how similar the two doxai are in degree, the former doxa is related to outright belief while the latter doxa is related to lack of outright belief—and so the two doxai are themselves qualitatively very different in crucial respects (to put it in more general terms relating to a contrast which has occupied philosophers for millennia, the challenge is to show why the inference to the conclusion that one's doxa in α_n is not reliable does not fallaciously take a negligible difference in the two doxai's *quantity* to indicate a negligible difference in their *quality*).

Secondly, and most importantly, as I've already hinted at, this whole issue is

[28]Moreover, the worry should be addressed that the required identifications may be unacceptable by almost everyone's lights, since, if one's doxa that one feels cold is > 0 at least in some cases in which one no longer feels cold, *transitivity of identity* would entail that, *even in* α_1, part of one's confidence that one feels cold is misplaced confidence, which may sceptically threaten one's knowledge that one feels cold in α_1.

[29]True, in setting up the argument, and even in the most recent discussion, I have myself attributed reliability and unreliability to doxai (confidences) rather than outright beliefs. But that has been merely for the sake of uniformity with the original presentation of the argument: it is only at this stage in the dialectic that it becomes important to see that a substantial, possibly crucial distinction needs to be made.

sidestepped by a more general and more powerful version of the argument just presented, which dispenses with outright beliefs and works only with doxai. Start by considering the property of degrees of doxa of being such that a mistaken doxa with that degree would give rise to a charge of unreliability against the knowledgeability of a doxa[30] had in similar enough cases (let's call such property 'the property of being *r-dangerous*'). By the properties of the reals and the finitude of $\{x : 1 \leq x \leq 1,000,000,000\}$, there is a least real number $\mathfrak{m}$ such that, for some n, in α_n one's doxa that one feels cold is $\mathfrak{m}$ and one's doxa is r-dangerous. Let $\mathfrak{r}$ be such an $\mathfrak{m}$. Next, reflect that, by the least-number principle, there is a last m such that in α_m one's doxa that one feels cold is at least $\mathfrak{r}$. Let r be such an m. Given (CONF), it is uncontroversial that $r < 1,000,000,000$. Hence, in α_{r+1} it is simply not the case that one's doxa is r-dangerous, and so, even if one's doxa in α_{r+1} were mistaken, it is simply not the case that that would give rise to a charge of unreliability against one's doxa in α_r that one feels cold.

Return now again to that part of the AL-argument which derives (KMAR) from more fundamental principles concerning knowledge and reliability. At the crucial junction, in the sub-argument for *reductio*, we are asked to assume that, for an arbitrary n, in α_n one knows (and so has a reasonably high doxa) that one feels cold, but that in α_{n+1} one does not feel cold. We are asked then to infer from the fact that, by (CONF), in α_{n+1} one has still a considerably high doxa—at worst *ever so slightly lower* than the one one has in α_n—that one's doxa in α_n is not reliable. The inference is supposed to be licenced by the consideration that a *very similar doxa* in a very similar condition formed on a very similar basis in a very similar case—i.e. the doxa one has in α_{n+1} in ⟨One feels cold⟩—is mistaken. But it should by now be clear that such a consideration concerning the high similarity in the doxai is irrelevant. For it is now open to the friend of (LUM) to retort that, for all the AL-argument has shown, n may well be r, in which case, although one's doxa in α_{n+1} is very similar to one's doxa in α_n, in α_{n+1}—as opposed to α_n—one's doxa is *no longer r-dangerous*. And if in α_{n+1} one's doxa is no longer r-dangerous, it is simply wrong to think that the fact that in α_{n+1} one does not feel cold would give rise to a charge of unreliability against the knowledgeability of one's doxa in α_n.

To compare with our guiding example of vagueness, the inference from an arbitrary n being such that in α_n one's doxa is r-dangerous to $n + 1$ also being such would be no less fallacious than the inference from an arbitrary n being such that a man with n hairs is bald to $n + 1$ also being such. The suggestion advanced at the end of section 4 is thus substantiated: *insofar as* (KMAR) is supposed to be justified by the fallacious derivation just discussed, it is indeed the case that this principle receives all of its plausibility in a spurious way from vagueness (in particular, from the sorites susceptibility of 'is a degree of doxa such that, in the situation described by the AL-argument, a doxa with that degree to the effect that one feels cold is r-dangerous'), and hence should be rejected on grounds that are quite independent from any issue concerning luminosity.

[30] For the sake of uniformity, I'm now attributing knowledgeability and non-knowledgeability to doxai rather than outright beliefs. I don't think that the difference is relevant in this context, and everything I'm saying, suitably recast, would apply just as well if we attributed knowledgeability and non-knowledgeability to outright beliefs.

There is more bad news for (KMAR). For not only does the previous argument show that its *justification* is radically flawed; it also provides the materials to question its *truth*. To see this, reflect that it is open whether:

(1) In α_r one knows that one feels cold

is true, and it is open whether:

(2) In α_{r+1} one does not feel cold

is true; indeed, unless one has antecedent reasons against (LUM), it would also seem to be open whether (1) and (2) are both true, and hence open whether (KMAR) has a false instance. Very interestingly, in such a situation, one's doxa would *stay reliable* in α_r, since, although it is the case that in α_{r+1} one does not feel cold, it is also the case that in α_{r+1} one's doxa drops enough (even if ever so slightly so) as to cease to be r-dangerous. In such a situation, then, one would in effect have *reliable knowledge at the limit*. This shows how, far from severing the connection between knowledge and reliability, the present argument against (KMAR) fully honours that connection and shows that (KMAR) is only spuriously justified by it.

Of course, a friend of (LUM) may go on and accept that (1) and (2) are both true (see section 8 for a plausible formal model of how doxai would look like in such a situation); a foe of (LUM) must reject this possibility. (This is so because, given the monotonicity fact that, for every n_0 and n_1 such that $n_0 \leq n_1$, if in α_{n_1} one knows that one feels cold, so does one in α_{n_0}, and, if in α_{n_0} one does not feel cold, neither does one in α_{n_1}, such possibility would in effect entail (LUM).) For those who don't have antecedent reasons for or against (LUM), pending further argument, recognition of this possibility as a possibility, and agnosticism about its realisation, are mandated (and so recognition of the possibility that (KMAR) is false as a possibility, and agnosticism about its realisation, are also mandated). As we've just seen, this attitude is perfectly consistent with the full recognition of the connection between knowledge and reliability.

7 Objections and Replies

The last considerations point to a possible line of defence on behalf of (KMAR).[31] Namely, (KMAR) could be saved from the argument of the previous section if it could be shown that, at least in the situation described by the AL-argument:

(A) Either one can only know that one feels cold if one's doxa that one feels cold is *higher than* $\mathfrak{r}$;

[31] I'm here indebted to discussions with Martin Smith.

(B) Or one can only fail to feel cold if one's doxa that one feels cold is *lower than the degree* $\mathfrak{r}'$ of one's doxa in α_{r+1} that one feels cold.

If (A) were correct, if one's doxa that one feels cold is $\mathfrak{r}$, it would follow that one does not after all know that one feels cold. But then, in the problematic case where n is r, the relevant instance of (KMAR) would have a false antecedent and so be true. If (B) were correct, if one's doxa that one feels cold is $\mathfrak{r}$, it would follow that one does after all still feel cold in the next case. But then, in the problematic case where n is r, the relevant instance of (KMAR) would have a true consequent and so be true.

(KMAR) saved? That will of course depend on whether either (A) or (B) can independently be supported. In this respect, (B) does not hold out much hope, given the speculative character of its postulation of a connection between one's feeling cold and the degree of one's doxa that one feels cold. Let's set it aside and focus on the more promising (A), which postulates rather a connection between one's knowing that one feels cold and the degree of one's doxa that one feels cold. While it is extremely plausible that there is at least *some kind or other* of connection between these two states, at first blush it too might seem to come out of the blue that the connection in question is of the *specific kind* postulated by (A). Indeed, quite the opposite would seem far more plausible: if one's doxa is high enough as to count as r-dangerous, it would also seem high enough as to count as knowledgeable if it also exhibits the appropriate epistemic properties. A highly non-trivial argument would be needed to subvert this natural intuition. (I will address an important source of the temptation to run such an argument in section 9, once a concrete formal model of the situation described by the AL-argument has been introduced, arguing that (A) is actually inconsistent with (FORM) and (CONF)).

I think that rather than seeking for a rather unlikely argument teasing apart the degree of doxa required for knowledgeability and the degree of doxa required for r-dangerousness, the best objection in the vicinity that can be raised against the newer version of the argument should be more indirect and go along the following lines. Let's assume the belief requirement on knowledge:

(KB) One knows that P only if one believes that P.

Then in α_r one can only know that one feels cold if in α_r one outright believes that one feels cold. Let's also assume that there is a minimal link between outright belief and r-dangerousness:

(OBRD) For every n, if in α_n one outright believes that one feels cold, then in α_n one's doxa that one feels cold is r-dangerous,

which should be uncontroversial. Then the advertised "more general and more powerful version of the argument" that I have been presenting requires in effect that it at least be open whether, in addition to marking a sharp boundary *for r-dangerousness*, α_r also marks a sharp boundary *for outright belief*. For, in order for it to be conclusive, this version of the argument requires that it be open whether (1) and (2) are both true. By (KB), (1) entails that in α_r one outright believes that one feels cold. Together with (2), it also entails, by (OBRD), that in α_{r+1} one does not outright believe that one feels cold. Since in this context we can safely assume that a consequence of something open is also open, this shows that this version of the argument requires that it be open whether α_r also marks a sharp boundary for outright belief. For all intents and purposes, this would seem to reduce this version of the argument to the previous, allegedly "simpler", one.

I am myself very sceptical as to whether this is the case, since, while (OBRD) seems a safe assumption in this context, (KB) is not, for the reasons adumbrated in fn 1. Notice, interestingly, that the AL-argument was careful enough as not to assume (KB), but only the possibly weaker principle that knowledge requires reasonably high doxa (confidence). (As I've mentioned in fn 6, I'd be a little sceptical about this latter principle as well, but, as I've already said there, I wish to set that issue aside in this paper.)

Nevertheless, grant (KB). Then it is correct to say that the newer version of the argument requires that it be open whether α_r also marks a sharp boundary for outright belief, and so, at least in a nominal sense, is "reduced" to the older version of the argument. Incidentally, this should of course not be such a great solace for the AL-argument, since the reader will recall that the upshot of our discussion of the older version was far from encouraging for the AL-argument. Be that as it may, the all-important point is that the supposed "reduction" overlooks the very different framework of the newer version, which focusses on doxai rather than outright beliefs. This framework helps to bring out how feeble the resistance to the older version is. Recall from the previous section that that resistance was forced staunchly to attribute reliability and unreliability to doxai rather than to outright beliefs. The newer version shows how, as soon as some basic formal properties of doxai's reliability are worked out in terms of r-dangerousness, the charge that the derivation of (KMAR) (and (KMAR) itself) ignores the sharp boundary for outright belief can straightforwardly be recast as the charge that the derivation of (KMAR) (and (KMAR) itself) ignores the sharp boundary for r-dangerousness: the mere switch from outright-belief reliability to doxa reliability does nothing to ameliorate the situation. Given (OBRD) and (KB), the newer version is indeed committed to its being open that the sharp boundary for r-dangerousness coincides with the sharp boundary for outright belief, but the high independent plausibility of this commitment under (KB)[32] only serves to highlight even more the unwarrantedness of the crucial inference in the

[32] In our dialectic, the crucial part of the open coincidence is constituted by its being open that the sharp boundary for r-dangerousness comes no later than the sharp boundary for outright belief (that, conversely, it comes no sooner is ensured by the uncontroversial (OBRD)). As has already been indicated in the text, it is highly plausibly open—quite independently of issues of luminosity—that a doxa high enough as to be r-dangerous is also high enough as to be sufficiently strong for knowledge. Given (KB), it will also be open that it is high enough as to determine outright belief, from which it straightforwardly follows that it is open that the sharp boundary for r-dangerousness comes no later than the sharp boundary for outright belief. (This way of motivating the commitment to the open coincidence between the sharp boundary for r-dangerousness and the sharp boundary for outright belief is independent from the way in which the argument in section 6 against (KMAR) is so committed because it does not require that it be open that, when one's doxa has reached the sharp boundary for r-dangerousness, one still knows that

derivation of (KMAR) (and the open possibility of (KMAR)'s falsity), as well as the shaky grounds of the idea that, although no longer high enough for outright belief, one's doxa in α_{n+1} is still high enough as to give rise to a charge of unreliability against the knowledgeability of one's doxa in α_n (in any event, as I've already said, I'll argue in section 9 that that idea is inconsistent with (FORM) and (CONF)).

A different line of defense on behalf of the AL-argument appeals to the possibility of running a *modal* rather than *temporal* version of the argument.[33] My argument has focussed on the impossibility of establishing that in α_r one's doxa that one feels cold is unreliable on the basis of the hypothetical fact that in α_{r+1} one's doxa that one feels cold is mistaken: I have argued that α_{r+1} does not impugn the reliability of one's doxa in α_r. This is fine as far as it goes, but it does not of course provide a guarantee that there is no case whatsoever that impugns the reliability of one's doxa in α_r. And, indeed, once we start looking at cases which, albeit metaphysically possible, are not part of the situation described by the AL-argument, candidates for impugning the reliability of one's doxa in α_r do seem to exist. The most promising kind of case is the metaphysically possible kind of case in which one's doxa is just the same as one's doxa in α_r but one just barely does not feel cold. Let's assume that there is a case β of this kind that is most similar to α_r. Uncontroversially, β is a metaphysically possible case, *quite similar* to α_r, in which one's doxa that one feels cold is both mistaken *and r-dangerous*. The crucial question however is whether β is *similar enough* to α_r, in the sense of 'similar' that is relevant to reliability (henceforth, let's use 'similar*' and its relatives to express such a sense). If it were, in α_r one's doxa that one feels cold would not be reliable, and hence one could not have reliable knowledge at the limit, contrary to what I've been defending.

I think that a fatal dilemma awaits any attempt to establish the antecedent of the previous conditional. Either the similarity* relation is *transitive* or it is *not*. Assume first that it is transitive. Then, since we can safely assume that it is also reflexive and symmetric, similar* cases form cells of a partition of the set of cases and, given standard logical and set-theoretical assumptions, such cells will be delimited by sharp boundaries. But if so, since there are chains of pairwise arbitrarily highly resembling cases whose extremes clearly lie in different similarity* cells (for example, a chain that starts with a case in which one is quite accurate in one's judgements about one's thermal sensations and ends with a case in which one has gone completely crazy and would judge to feel cold even while feeling extremely hot), there are bound to be pairs of cases γ_0 and γ_1 that lie in different similarity* cells although they resemble one another in an arbitrarily high fashion. And, for all the AL-argument and the latest objection have shown, γ_0 and γ_1 may well be α_r and β respectively.

Assume second that the similarity* relation is not transitive. How could we then establish that β is similar* to α_r? Well, we could *now* plausibly assume:

one feels cold. It only requires that it be open that, when one's doxa has reached the sharp boundary for r-dangerousness, one's doxa is still sufficiently strong for knowledge, but it is compatible with its being ruled out that, in such a case, one still knows that one feels cold—say, because one's doxa fails to meet an epistemic condition on its being knowledgeable.)

[33] I'm here indebted to discussions with Julien Dutant.

(I) If in δ_0 one's doxa that one feels cold is $\mathfrak{d}$ and one feels cold to degree $\mathfrak{f}$, and in δ_1, for some $\mathfrak{e} \leq \mathfrak{b}$, one's doxa that one feels cold is $\mathfrak{d} + \mathfrak{e}$ (alternatively, $\mathfrak{d} - \mathfrak{e}$) and one feels cold to degree $\mathfrak{f} + \mathfrak{c}\mathfrak{e}$ (alternatively, $\mathfrak{f} - \mathfrak{c}\mathfrak{e}$), then δ_1 is similar* to δ_0;

(II) If in δ_0 one's doxa that one feels cold is $\mathfrak{d}$ and one feels cold to degree $\mathfrak{f}$, and in δ_1, for some $\mathfrak{e} \leq \mathfrak{b}$, one's doxa that one feels cold is $\mathfrak{d} \pm \mathfrak{e}$ and one still feels cold to degree $\mathfrak{f}$, then δ_1 is similar* to δ_0,

where $\mathfrak{b}$ is a suitable bound and $\mathfrak{c}$ a suitable coefficient. We could then assume that there is a suitable choice of $\mathfrak{b}$ and $\mathfrak{c}$ such that (I) entails that α_{r+1} is similar* to α_r (as, of course, we've been assuming throughout) and (II) entails that β is similar* to α_{r+1}. However, under the assumption of non-transitivity of the similarity* relation, that α_{r+1} is similar* to α_r and β is similar* to α_{r+1} does no longer suffice to establish that β is similar* to α_r. It is at best very unclear how else we could establish that β is similar* to α_r. And it is at best very unclear that simply insisting that it is just obvious that β is similar* to α_r—in effect, that at some time in the situation described by the AL-argument one's confidence might easily have been mistaken—amounts to anything more interesting than simply insisting that one is not an optimal subject.

Rather than simply insisting that it is just obvious that β is similar* to α_r, one could appeal to:

(III) If in δ_0 one's doxa that one feels cold is $\mathfrak{d}$ and one feels cold to degree $\mathfrak{f}$, and in δ_1, for some $\mathfrak{e} \leq \mathfrak{b}$, one feels cold to degree $\mathfrak{f} \pm \mathfrak{e}$ and one's doxa that one feels cold still is $\mathfrak{d}$, then δ_1 is similar* to δ_0,

where, again, $\mathfrak{b}$ is a suitable bound. (III) would do the required job, as, given our assumptions about α_r and β, it would entail that β is similar* to α_r. But I don't think that it brings us much further than the simple insistence just discounted. Notice that (III) is in some sense the converse of (II): while (II) says, roughly, that small oscillations in doxa with a fixed thermal sensation are sufficient for similarity*, (III) says, roughly, that small oscillations in thermal sensation with a fixed doxa are sufficient for similarity*. But given that small oscillations in thermal sensation are sometimes oscillations from feeling to not feeling cold, (III) in fact says that, in some cases, one is completely insensitive as to whether one feels cold or not. Hardly an uncontroversial assumption![34]

[34] A completely different concern—albeit not one that everyone will be sensitive to—may be constituted by the use of *classical* logic made by my argument against (KMAR) in reasoning about *vague* matters. However, even setting aside that such concern worryingly ties together the AL-argument and vagueness (see sections 3 and 4), it's easy to see that, on the many non-classical logics of vagueness *rejecting* sorites premises, the derivation of (KMAR) is just as fallacious as before, and that, on the few non-classical logics of vagueness *accepting* sorites premises (see Zardini [2008a]; Zardini [2008b], pp. 93–173; Zardini [2009]; Zardini [2012b]), even if the AL-argument may be compelling up to (SOR), the latter is now a harmless conclusion perfectly consistent with (POSNEG). As for the more general point that reliable knowledge *in the penumbra* is possible, once classical logic is weakened we can obviously no longer claim that reliable knowledge *at the limit with falsity* is possible, for, by factivity of knowledge, that would entail a sharp

8 A Formal Model of Reliable Luminosity

It will prove helpful to provide broadly probabilistic models of the situation of reliable knowledge at the limit whose possibility has been defended in section 6. In the light of the discussion in section 5, it will be more realistic to use as underlying measure functions of doxa a class of super-additive functions, and, for purely illustrative purposes (see fn 19), I will focus on Dempster-Shafer functions (henceforth, as usual and—in the light of the discussion in section 5—appropriate, 'belief functions').

Definition 4. Given a finite set U, a *basic belief assignment* (bba) on U is a function $\mathfrak{M} : \wp(U) \mapsto \mathbb{R}$ such that:

(i) $\mathfrak{M}(\varnothing) = 0$;

(ii) $\sum_{X \subseteq U} (\mathfrak{M}(X)) = 1$.

In our context, a bba represents the doxastic investment a subject makes in a hypothesis considered *in its specificity* (i.e. not because of its being a weaker consequence of a hypothesis on which a certain doxastic investment has been made). Belief functions can be defined in terms of bbas as follows:

Definition 5. Given a finite set U and a bba $\mathfrak{M}$ on U, a *belief function* on U is a function $\mathfrak{B} : \wp(U) \mapsto \mathbb{R}$ such that:

(BF) $\mathfrak{B}(X) = \sum_{Y \subseteq X} (\mathfrak{M}(Y))$.

The following fundamental properties of belief functions are then easily derivable from (BF):

Theorem 1. *Suppose that $\mathfrak{B}$ is a belief function on U. Then:*

(i) $\mathfrak{B}(\varnothing) = 0$*;*

(ii) $\mathfrak{B}(U) = 1$*;*

(iii) If $X \subseteq Y$, then $\mathfrak{B}(X) \leq \mathfrak{B}(Y)$;

(iv) For every $n \in \mathbb{N}^*$, for every $X_1, X_2, X_3 \ldots, X_n \subseteq U$,
$$\mathfrak{B}\left(\bigcup_{1 \leq m \leq n} (X_m)\right) \geq \sum_{\varnothing \neq I \subseteq \{1,2,3\ldots,n\}} \left((-1)^{|I|+1}\right)\mathfrak{B}\left(\bigcap_{i \in I} (X_i)\right).$$

boundary for one's feeling cold, avoiding the commitment to which was the whole point of the exercise of weakening classical logic. What we can still claim on the strength of the argument against (KMAR) is that, given a certain non-classical theory of vagueness, for every n_0 such that the theory does not rule out α_{n_0} as being a case in which one feels cold, reliable knowledge in α_{n_0} that one feels cold is possible, no matter how close an n_1 exists such that the theory does rule out α_{n_1} as being a case in which one feels cold—one can reliably know that one feels cold even in the last case*s* in which one feels cold.

Note that, if X and Y are incompatible (i.e. $X \cap Y = \varnothing$), property (iv) yields as a special case the *super-additivity* property $\mathfrak{B}(X \vee Y) \geq \mathfrak{B}(X) + \mathfrak{B}(Y)$.

To fix ideas, let's assume that the set of cases corresponding to ⟨One feels cold⟩ is $\{\alpha_n : 1 \leq n \leq 500,000,000\}$, and that the threshold for a doxa's r-dangerousness is $.000000002 = \mathfrak{t}$. We can then provide the following formal model of the dynamics of doxai in the situation described by the AL-argument, a model which corresponds to the possibility of reliable knowledge at the limit defended in section 6 and whose coherence and plausibility show both that the derivation of (KMAR) is fallacious and that it is open that (KMAR) itself is false:

Definition 6. Let α_1, α_2, $\alpha_3 \ldots$, $\alpha_{1,000,000,000}$ be as before. Let $\mathscr{L} = \langle A, B \rangle$ be as follows:

- $A = \{\alpha_n : 1 \leq n \leq 1,000,000,000\}$;
- $B = \{\mathfrak{B}_n : 1 \leq n \leq 1,000,000,000\}$, where $\mathfrak{B}_n$ is the belief function generated by the bba $\mathfrak{M}_n$ such that, for every $n \leq 500,000,000$:
  - $\mathfrak{M}_n(\{\alpha_m : 1 \leq m \leq 500,000,000\}) = 1,000,000,000 - 2(n-1)/1,000,000,000$;
  - $\mathfrak{M}_n(A) = 1 - \mathfrak{M}_n(\{\alpha_m : 1 \leq m \leq 500,000,000\})$;
  - For every other X, $\mathfrak{M}_n(X) = 0$,

  and such that, for every $n \geq 500,000,001$:
  - $\mathfrak{M}_n(\{\alpha_m : 500,000,001 \leq m \leq 1,000,000,000\}) = 1,000,000,000 - 2(1,000,000,000 - n)/1,000,000,000$;
  - $\mathfrak{M}_n(A) = 1 - \mathfrak{M}_n(\{\alpha_m : 500,000,001 \leq m \leq 1,000,000,000\})$;
  - For every other X, $\mathfrak{M}_n(X) = 0$.

Finally, let's make some simplifying but plausibly not distorting assumptions about the connections between knowledge and r-dangerousness. One knowledge-threatening factor at issue in the AL-argument is the degree of doxa. In light of the previous discussion, we can set the degree of doxa necessary for knowledge equal with the threshold for r-dangerousness. Since the only other knowledge-threatening factor at issue in the AL-argument is reliability, we can safely if artificially assume that a doxa is knowledgeable *iff* it is r-dangerous and reliable. In turn, for our purposes we can safely take it that a doxa is reliable iff there are no very similar cases in which a very similar *and still r-dangerous*[35] doxa in a very similar condition formed on a very similar basis is false.

From the earlier assumptions about what set of cases corresponds to ⟨One feels cold⟩ and what the threshold for a doxa's r-dangerousness is, from the later assumptions about the dynamics of doxai and from the latest assumptions about the connections between knowledge and r-dangerousness, we are at last in a position to draw inferences about

[35]This crucial qualification is of course one of the main lessons of the discussion of section 6.

knowledge of ⟨One feels cold⟩ in the model $\mathscr{L}$. Under all these assumptions, it is very interesting to observe that a *neighbourhood* model of epistemic logic is actually already implicit in $\mathscr{L}$ (see Mongin [1994] for similar but more general results).

Definition 7. For every n, the *essential pre-image* of $\mathfrak{B}_n$ is $\mathsf{e}(\mathfrak{B}_n) = \{X : \mathfrak{B}_n(X) \geq \mathfrak{t}\}$.

Theorem 2. *For every n:*

(M′) $\mathsf{e}(\mathfrak{B}_n)$ *is* supplemented *(i.e. for every X and $Y \subseteq A$, if $X \in \mathsf{e}(\mathfrak{B}_n)$ and $X \subseteq Y$, then $Y \in \mathsf{e}(\mathfrak{B}_n)$);*

(C′) $\mathsf{e}(\mathfrak{B}_n)$ *is* closed under finite intersections *(i.e. for every X and $Y \subseteq A$, if $X \in \mathsf{e}(\mathfrak{B}_n)$ and $Y \in \mathsf{e}(\mathfrak{B}_n)$, then $X \cap Y \in \mathsf{e}(\mathfrak{B}_n)$);*

(N′) $\mathsf{e}(\mathfrak{B}_n)$ contains the unit *(i.e. $A \in \mathsf{e}(\mathfrak{B}_n)$);*

(D′) $\mathsf{e}(\mathfrak{B}_n)$ *is* proper *(i.e. for every $X \subseteq A$, if $X \in \mathsf{e}(\mathfrak{B}_n)$, then $\overline{X} \notin \mathsf{e}(\mathfrak{B}_n)$);*

(T′) $\mathsf{e}(\mathfrak{B}_n)$ *is* reflexive *(i.e. for every $X \in \mathsf{e}(\mathfrak{B}_n)$, $\alpha_n \in X$);*

(4′) $\mathsf{e}(\mathfrak{B}_n)$ *is* transitive *(i.e. for every $X \in \mathsf{e}(\mathfrak{B}_n)$, $\{\alpha_m : X \in \mathsf{e}(\mathfrak{B}_m)\} \in \mathsf{e}(\mathfrak{B}_n)$).*

Theorem 2 in effect ensures that essential pre-images can be used as (very well-behaved) local neighbourhood functions in order to generate from $\mathscr{L}$ the following neighbourhood model of epistemic logic:

Definition 8. Take a standard sentential mono-modal language with K as its set of sentential constants and $\mathcal{K}$ as its modal operator. Let $\mathscr{L}^* = \langle A, N, \mathsf{v}\rangle$ be the following neighbourhood model:

- A is as before;
- $N : A \mapsto \wp(\wp(A))$ is such that, for every n, $N(\alpha_n) = \mathsf{e}(\mathfrak{B}_n)$;
- $\mathsf{v} : K \mapsto \wp(A)$ is such that $\mathsf{v}(P_0) = \{\alpha_n : 1 \leq n \leq 500,000,000\}$.

Thus, in our modelling P_0 expresses ⟨One feels cold⟩. v can recursively be extended to a complete interpretation function i as usual, where in particular the clause for $\mathcal{K}$ is as follows:

(K) $\mathsf{i}(\mathcal{K}\varphi) = \{\alpha_n : \mathsf{i}(\varphi) \in N(\alpha_n)\}$.

We now have to check that the coherent and plausible model $\mathscr{L}^*$ is indeed a model which corresponds to the possibility of reliable knowledge at the limit defended in section 6, and in particular that it is a model where ⟨One feels cold⟩ is luminous, knowledge is reliable and (KMAR) fails. That $\mathscr{L}^*$ has these three properties is ensured by the three following corresponding theorems (writing $\alpha_n \models_{\mathscr{L}^*} \varphi$ for $\alpha_n \in \mathsf{i}(\varphi)$):

Theorem 3. *For every* n, $\alpha_n \models_{\mathscr{L}^*} P_0 \supset \mathcal{K}P_0$.

Theorem 4. *For no* n *such that* $\mathfrak{B}_n(P_0) \geq \mathsf{t}$, $\alpha_n \models_{\mathscr{L}^*} \neg P_0$.

Theorem 5. $\alpha_{500{,}000{,}000} \models_{\mathscr{L}^*} \mathcal{K}P_0$ *and* $\alpha_{500{,}000{,}001} \models_{\mathscr{L}^*} \neg P_0$.

We can also check that such modelling does not force us to give up on any interesting formal property that one might wish to associate with the logic of knowledge. The result is already semantically encapsulated in theorem 2 and is made explicit in the following theorem:

Theorem 6. *For every* n*:*

(M) $\alpha_n \models_{\mathscr{L}^*} \mathcal{K}(\varphi \wedge \psi) \supset \mathcal{K}\varphi \wedge \mathcal{K}\psi$;

(C) $\alpha_n \models_{\mathscr{L}^*} \mathcal{K}\varphi \wedge \mathcal{K}\psi \supset \mathcal{K}(\varphi \wedge \psi)$;

(N) $\alpha_n \models_{\mathscr{L}^*} \mathcal{K}\top$;

(D) $\alpha_n \models_{\mathscr{L}^*} \mathcal{K}\varphi \supset \neg\mathcal{K}\neg\varphi$;

(T) $\alpha_n \models_{\mathscr{L}^*} \mathcal{K}\varphi \supset \varphi$;

(4) $\alpha_n \models_{\mathscr{L}^*} \mathcal{K}\varphi \supset \mathcal{K}\mathcal{K}\varphi$.

Of course, for many of these properties, it is a matter of high controversy in epistemology whether knowledge has them—to repeat, the point is rather that our modelling does not force us to give up on any of them. An opposite worry would be that our modelling may force us to accept all such properties (as holding at least in the situation described by the AL-argument, at least when that situation exhibits reliable luminosity), but that worry is easily addressed by observing that similar models can be given that are less constrained than $\mathscr{L}^*$ and that thus fail to verify some or all of them.[36]

[36] I've used the framework of *neighbourhood* semantics for modal logics—rather than the philosophically more familiar framework of *relational* semantics—for two reasons. Firstly, because, in its systematic discriminating power with respect to logics that fail any of (M), (C) and (N), neighbourhood semantics is widely recognised as a better and more conspicuous tool for doing *epistemic* logic than standard relational semantics is. Secondly, because neighbourhood frames and models can be generated out of belief functions in the way exploited in theorem 2. That being said, I hasten to add that, given that e enjoys certain crucial properties, our neighbourhood model can easily be converted into a relational model. Let's first observe that, in addition to (C$'$), e also exhibits the following infinitary strengthening of that property:

(C$'^{+}$) $\mathsf{e}(\mathfrak{B}_n)$ is *closed under arbitrary intersections* (i.e. for every $F \subseteq \wp\wp(A)$, if, for every $X \in F$, $X \in \mathsf{e}(\mathfrak{B}_n)$, then $\bigcap(F) \in \mathsf{e}(\mathfrak{B}_n)$).

(notice that (N$'$) is a special case of (C$'^{+}$) got by setting $F = \varnothing$). Given that e enjoys both (M$'$) and (C$'^{+}$), we can then convert our neighbourhood model $\mathscr{L}^*$ into a relational model $\mathscr{L}^{**} = \langle A, R, \mathsf{v}\rangle$ by setting $[\langle\alpha_n, \alpha_m\rangle \in R$ iff $\alpha_m \in \bigcap(N(\alpha_n))]$ (the basic ideas behind neighbourhood semantics for modal logics trace back as far as McKinsey and Tarski [1944] and have first been developed in the pioneering works of Montague [1968] and Scott [1970]).

9 Confidence Requirements on Knowledge

With a concrete formal model in place, I would like to go back to an objection raised in section 7 and address an important source of the temptation of arguing in favour of (A). Looking at $\mathscr{L}^*$ and at its identification of $\mathfrak{r}$ with $\mathfrak{t} = .000000002$, one might protest that surely a doxa that low cannot qualify as knowledge, and hence that (A) holds and (KMAR) is not threatened at $\alpha_{500,000,000}$.

Before offering my main reply to what I take to be the spirit of such protest, I'd like first to clear the ground of a likely misunderstanding. The protest might be misled by the use of a non-classical probability calculus, which, while sharing with the classical calculus the normalisation of the probability measure to the real interval $[0, 1]$, in effect models (CONF), as it were, "*in sequence*", by first gradually decreasing one's doxa that one feels cold and *then* gradually increasing one's doxa that one does not feel cold. This accords well with the points made in section 5, centring around the apparent fact that one can rationally diminish one's doxastic investment in $\langle C \rangle$ without thereby augmenting one's doxastic investment in ⟨It is not the case that C⟩. On the first half of the way from dawn to noon, one gradually lowers one's doxa that one feels cold in reaction to the poorer and poorer evidence to that effect, without thereby raising one's doxa that one does not feel cold. This is quite understandably so, since, on that first half, one is not gaining any evidence to the effect that one does not feel cold—the thermal sensations are indeed less and less of cold, but they are still overall sensations of cold rather than of non-cold (after all, one still feels cold rather than not), and it would seem perverse to raise one's doxa that one does not feel cold in reaction to sensations of cold. Had we adopted the classical probability calculus, we would have modelled (CONF), as it were, "*in parallel*", by gradually decreasing one's doxa that one feels cold and *at the same time* gradually increasing one's doxa that one does not feel cold. On this alternative picture, which I have just given reason to reject, the most natural modelling assumptions would give us in $\alpha_{500,000,000}$ a credence of degree .500000001. That is what, on the picture I have been developing in the previous section, a doxa of degree .000000002 should be thought as corresponding to (at least in the situation described by the AL-argument), and so, to the extent to which it is plausible that the former is sufficiently strong for knowledge, it will also be plausible that the latter is.[37]

[37] This is probably the best place also to remark that, if one so wishes, the construction in the previous section can easily be modified in order to incorporate into it an NFMU-theory of vagueness (see fn 17). Setting aside further complications required by a treatment of higher-order vagueness, the key idea would be, at the semantic level, to introduce an area of borderline cases stretching from α_i to α_j where the truth value of P_0 is indeterminate, and, at the psychological level, to let, for every $i \leq n \leq j$, $\mathfrak{B}_n(\{\alpha_m : 1 \leq m < i\}) = 0$ as well as $\mathfrak{B}_n(\{\alpha_m : j < m \leq 1,000,000,000\}) = 0$ (assuming that, oversimplifying a bit, we model the condition expressed by φ with the set of cases at which φ is determinately true; although I'll save the details to the reader, it is possible to modify the construction in order to allow for a more fine-grained modelling). This would model quite well the fact that, on an NFMU-theory, one should presumably withdraw all doxastic investment from borderline conditions (as there is no fact of the matter about them). Since we can safely assume that the relevant conditions are borderline only if their negations are, this requires a violation of the classical probability calculus for doxai (even while keeping classical logic fixed), and belief functions are an apt substitute (as noted by an NFMU-theorist like Field

Somewhat incidentally, but very importantly, the non-classical "in-sequence" modelling of (CONF), contrary to the classical "in-parallel" modelling, allows a conclusive dismissal by everyone's lights of the worry that one's doxa that one feels cold is not reliable in $\alpha_{500,000,000}$. For, on the specific "in-sequence" modelling of (CONF) offered in the previous section, in the next case $\alpha_{500,000,001}$ *one does not have any positive degree of doxa at all that one feels cold*, and that obviously should not give rise to any charge of unreliability against the knowledgeability of one's doxa in $\alpha_{500,000,000}$—more strongly, that should be considered as a *guarantee* that one's doxa that one feels cold in $\alpha_{500,000,000}$ and in all the preceding cases is reliable (at least as long as reliability only depends on what is going on between α_1 and $\alpha_{1,000,000,000}$). This is to be contrasted with any plausible "in-parallel" modelling, on which in $\alpha_{500,000,001}$ one does still have a positive degree of credence that one feels cold (in the specific "in-parallel" modelling sketched in the previous paragraph, the most natural modelling assumptions would give us in $\alpha_{500,000,001}$ a credence of degree .499999999).[38] We have thus hit another crucial bind for the AL-argument: as observed in section 5, the argument is only plausible if it focusses on doxai rather than credences, but, contrary to credences, doxai have independently motivated formal properties (i.e. super-additivity) that allow a straightforward modelling which conclusively establishes by everyone's lights the reliability of one's doxa that one feels cold throughout the situation described by the AL-argument.

Even with these clarifications in play, the protest might continue to attack the knowledgeability of the doxa in question (and so the knowledgeability of the corresponding credence of degree .500000001) on the alleged grounds that it is still too low. I've already mentioned several times, in sections 1 (in particular, fn 1), 2 (in particular, fn 6) and 7, grounds for being suspicious with such insistence, but, again, let that pass. One might have thought that the insistence can easily be satisfied by changing the modelling assumptions, in such a way as to have a higher doxa in $\alpha_{500,000,000}$, keeping the latter fixed as the sharp boundary for one's feeling cold. However, no very high doxa can be had in $\alpha_{500,000,000}$, keeping also fixed $\mathfrak{t}$ as the sharp boundary for r-dangerousness. This is so because in $\alpha_{500,000,000}$ one has to know that one feels cold (if (LUM) has to be preserved), but that requires one's doxa to be under $\mathfrak{t}$ in $\alpha_{500,000,001}$, which together with a very high doxa in $\alpha_{500,000,000}$ would create a drop in doxa inconsistent with (CONF). The protest finds then its last refuge in holding that no assignment of degree of doxa in $\alpha_{500,000,000}$ can be high enough as to satisfy an alleged confidence requirement on knowledge (involving a suitably high threshold) and at the same time low enough as to be consistent with both (CONF) and a reliability requirement on knowledge (involving a suitably low threshold for r-dangerousness).

Setting aside the legitimate question of what grounds there could be for holding the last claim (which, as I've mentioned in section 7, is deeply unobvious in its commitment

[2000]).

[38] Although I can't see why one would want to, I should note that one can of course give an "in-parallel" modelling also for doxai—in fact even one with exactly the same features as those contemplated in the text, since classical probability functions are just a particular (limit) case of belief functions. Having noted this, I'll henceforth assume that our modelling for doxai is something along the lines of the "in-sequence" modelling offered in the previous section.

to the threshold placed by the alleged confidence requirement being suitably higher than the threshold placed by the reliability requirement), it will be sufficient to point out that it constitutes a dialectical harakiri. For, assuming (as we've been doing) that the sharp boundary for one's feeling cold is modelled as being in $\alpha_{500,000,000}$, a modelling of (CONF) is only going to be acceptable to a friend of (LUM) insofar as it also places in $\alpha_{500,000,000}$ the sharp boundary for one's doxa (alternatively, one's credence) being to at least some degree in favour of the hypothesis that one feels cold (especially if it's assumed, very plausibly, that, if ⟨One feels cold⟩ is luminous, so is ⟨One does not feel cold⟩). Given that (CONF) requires *graduality* in this latter transition, this *already by itself suffices* to force one's doxa to be (more or less) .000000002 in $\alpha_{500,000,000}$ (alternatively, force one's credence to be (more or less) .500000001 in $\alpha_{500,000,000}$). One could then apply to this obvious fact whatever argument one has for rejecting that a .000000002 doxa (alternatively, a .500000001 credence) is sufficiently strong for knowledge and—hey presto!—one could already conclude that $\alpha_{500,000,000}$ is a case in which one feels colds but is not in a position to know that one feels cold—that is, a case in which (LUM) fails. Whatever the (dubious) merits of this new argument against (LUM) are, it is evident that it wholly and only hinges on the alleged confidence requirement on knowledge *and does nowhere depend on any reliability requirement*. It is thus completely different from the AL-argument, in both letter and spirit: while the AL-argument is an interesting argument to the effect that, for some n, in α_n one's doxa fails to meet an *epistemic* condition on its being knowledgeable (i.e. reliability), the new argument is a boring argument to the effect that, for some n, in α_n one's doxa fails to meet an alleged *purely psychological* condition on its being knowledgeable (i.e. strength). If (KMAR) can only be defended by appealing to the alleged confidence requirement on knowledge, it can only be defended by appealing to a principle that, given (CONF) and some other uncontroversial assumptions, would already yield a conclusive argument against (LUM) independent of any consideration having to do with reliability or any other epistemic property. If that is the only defense of (KMAR), then this crucial premise of the AL-argument can be justified only at the cost of making that whole style of argument redundant.

Would we at least have sacrificed the original AL-argument in favour of a more compelling argument? Not at all. Under the assumption of the alleged confidence requirement on knowledge, (CONF) would really amount, among other things, to a declaration that, at least in the central cases, one is not in effect taking a view strong enough for knowledge. Given the other assumption that one feels cold in the central case $\alpha_{500,000,000}$, (CONF) would thus be straightforwardly inconsistent with (FORM) (as has been pointed out already in section 1). (FORM)—just as well as (CONF)—is however as necessary to the new argument as it was to the AL-argument: for the new argument at best only directly establishes that, because of too low a doxa, one does not *know*, and only (FORM) allows the transition from that to one's not *being in a position to know*, thereby making the argument engage with (LUM) (rather than simply with (LUM^+)). The new argument thus requires both (CONF) and (FORM), plus an additional alleged confidence requirement on knowledge that makes those two assumptions jointly inconsistent. Far from committing a subtle fallacy of vagueness, the new argument—and any argument of the same ilk appealing to the alleged confidence requirement on knowledge—is simply incoherent.

Once this incoherence is spotted, it becomes easy also to spot the incoherence attaching to using (A) in this context (notice that the points made in the previous two paragraphs do not immediately apply to (A), for (A) is only a confidence requirement on knowledge *relative* to the sharp boundary for r-dangerousness, rather than an *absolute* confidence requirement on knowledge, which has been the target of the previous two paragraphs). To be effective in this context, (A) must really be interpreted as saying not only that the sharp boundary for r-dangerousness is *to some extent or other* lower than the sharp boundary for a doxa being sufficiently strong for knowledge, but also that the distance between these boundaries is *large enough* for (CONF) to imply that, in the situation described by the AL-argument, one's doxa is at some time still high enough as to be r-dangerous although no longer high enough as to be sufficiently strong for knowledge. However, given the other assumption that one's doxa that one feels cold stops being r-dangerous in the central case $\alpha_{500,000,001}$, (A) so interpreted together with (CONF) would imply that one's doxa that one feels cold is not high enough as to be sufficiently strong for knowledge in the central case $\alpha_{500,000,000}$. And, again, given the assumption that one feels cold in the central case $\alpha_{500,000,000}$, (A) so interpreted together with (CONF) would thus be straightforwardly inconsistent with (FORM). We thus see that the apparently innocuous (FORM)—which, as far as I know, is never critically discussed and hardly ever made explicit in the extant literature—has a damning strength that makes it inconsistent with (CONF) and some confidence requirements on knowledge (like an absolute requirement or (A)) that might be behind the felt plausibility of (KMAR).[39]

[39] Berker [2008] and Ramachandran [2009] also focus on the problem that, in the derivation of (KMAR), it cannot be assumed that in α_{n+1} one still outright believes that one feels cold (with Leitgeb [2002], pp. 200–205 being an earlier insightful discussion touching on very similar points). Since, as I've mentioned in section 6, that problem is explicitly raised by Williamson himself, who in fact argues that the possible lack of outright belief in α_{n+1} does not essentially affect the point about reliability on which the derivation of (KMAR) rests, that is only the beginning of the dialectic, and my treatment of it differs greatly from those of these authors, even focussing on its core and setting aside my approach to the role of confidence requirements on knowledge in the AL-argument and to the modal version of the argument (issues that are neglected by these other authors). They want to argue, as I do, that the fact that in α_{n+1} one has a certain degree of mistaken confidence that one feels cold need not give rise to a charge of unreliability against the knowledgeability of one's outright belief in α_n. (They do not notice that, when confidence is credence, straightforward counterexamples of the kind discussed in section 5 can be given to support their claim, and do not thematise the distinction between credence and doxa which, on the one hand, allows the derivation of (KMAR) to survive that kind of counterexample, but, on the other hand, also allows a straightforward modelling which conclusively establishes by everyone's lights the reliability of one's doxa that one feels cold throughout the situation described by the AL-argument.) However, they fail to consider the all-important property of r-dangerousness, and hence fail to see the completely general fact that, no matter what details are given of what makes a doxa reliable or not, the derivation of (KMAR) is bound to stumble over the sharp boundary for r-dangerousness. It is this completely general fact that allows one to see that not only is the particular derivation of (KMAR) fallacious, but also that there simply cannot be an issue of reliability concerning the knowledgeability of one's doxa in α_n when α_n is α_r (as long as reliability only depends on what is going on between α_1 and $\alpha_{1,000,000,000}$; I have then argued in section 7 that the point in fact extends to cover also what is going on in non-actual but similar cases). And it is this guarantee of reliability that makes it possible to have reliable knowledge at the limit—contrary to these other authors, my treatment *shows that such knowledge is possible* (as far as reliability considerations are concerned), rather than simply *blocking a particular argument to the effect that it is not* (the contrast here being the same as that between a consistency proof for a theory

10 The Possibility of Reliable Knowledge at the Limit

I have argued that proper consideration of the extent to which vagueness does not imply the absence of sharp boundaries reveals the possibility of reliable knowledge at the penumbral limit with falsity, in the sense that reliability considerations do not preclude the existence of knowledge at the penumbral limit with falsity. I have not argued that such knowledge is ever realised, and in particular have not argued that no other kinds of considerations preclude its existence, although I should record my sympathy for the view that such knowledge is indeed realised at least in the situation described by the AL-argument (see fn 40). Leaving the discussion of those further matters for future work, I then conclude this one by noting that, at least as far as the issues I have discussed are concerned, careful heeding of vagueness and its phenomena (in particular, sorites susceptibility), far from forcing new and surprising limits on our knowledge, actually removes one of the main barriers—unreliability—often thought to stand in its way.[40]

---

and the identification of a particular fallacious derivation of a contradiction from the theory; the formal model in the previous section constitutes in effect such a proof). Let me exemplify this point with some more specific comments on Berker [2008], which is the most elaborated alternative defence of the idea that the fact that in α_{n+1} one has a certain degree of mistaken confidence that one feels cold need not give rise to a charge of unreliability against the knowledgeability of one's outright belief in α_n. Berker argues that a certain reliability constraint on confidence does not work for a subject whose confidence in the situation described by the AL-argument changes continuously, apparently for the only reason that if it did, luminosity would fail for any such subject. He is thereby assuming that it is possible for a continuous confidence to be reliable throughout the situation described by the AL-argument—given that the argument's only assumption about confidence is (CONF), which is entailed by a confidence's continuity, and given that the argument's target is precisely to show that any such confidence cannot be reliable throughout the situation described by the argument, it is hard to see how Berker's assumption can serve as a basis for an effective criticism of the argument. Moreover, even granting that such criticism were effective, it would remain just that—a criticism of [a particular argument to the effect that reliable knowledge at the limit is not possible]. As such, it leaves wide open that other reliability constraints might do the job in the AL-argument. And, more importantly, it leaves unscathed the AL-argument as applied to any subject whose confidence in the situation described by the argument does not change continuously; furthermore, since it would seem that merely being continuous rather than non-continuous does not afford epistemic advantages in the situation described by the AL-argument, it would seem that, if luminosity fails for any subject with a non-continuous confidence, it will after all fail also for any subject with a continuous confidence.

[40]Suppose, contrary to NFMU-theories of vagueness, that, for some α, although in α it is borderline whether one feels cold, in α it is a fact of the matter that one feels cold. Couldn't we then bypass the sophistications of the AL-argument and straightforwardly point out that (LUM) fails at α because, being borderline in α whether one feels cold, in α one cannot know that one feels cold, although in α one does feel cold? Not so quick. If it's really a fact of the matter in α that one feels cold, then in α one presumably has a *conclusive ground* to believe that one feels cold—namely, one's feeling cold. Moreover, that in α one is in a position to know that one feels cold only requires that, in α, *if* one came to believe that one feels cold on that ground, one *would* come to know that one feels cold. Under these assumptions, it's not at all clear that that counterfactual condition is not met by one in α, and so not at all clear that (LUM) fails at α. In fact, I've argued in this paper against the usual allegation to the effect that the counterfactual condition is not met by one in α because one's doxa would not be reliable, so that at least that reason for maintaining that the counterfactual condition is not met by one in α is now off the

References

Selim Berker. Luminosity regained. *Philosophers' Imprint*, 8:1–22, 2008.

Earl Conee. The comforts of home. *Philosophy and Phenomenological Research*, 70: 444–451, 2005.

Arthur Dempster. Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics*, 38:325–339, 1967.

Didier Dubois and Henri Prade. *Théorie des possibilités*. Masson, Paris, 1985.

Hartry Field. Indeterminacy, degree of belief, and excluded middle. *Nous*, 34:1–30, 2000.

Kit Fine. Vagueness, truth and logic. *Synthese*, 30:265–300, 1975.

Hannes Leitgeb. Review of Timothy Williamson, *Knowledge and its Limits*. *Grazer Philosophische Studien*, 65:195–205, 2002.

John McKinsey and Alfred Tarski. The algebra of topology. *Annals of Mathematics*, 45: 141–191, 1944.

Philippe Mongin. Some connections between epistemic logic and the theory of nonadditive probability. In Paul Humphreys, editor, *Patrick Suppes: Scientific Philosopher*, volume I: Probability and Probabilistic Causality, pages 135–171. Kluwer, Dordrecht, 1994.

Richard Montague. Pragmatics. In Raymond Klibansky, editor, *Contemporary Philosophy: A Survey*, volume I: Logic and Foundations of Mathematics, pages 102–122. La Nuova Italia Editrice, Florence, 1968.

Colin Radford. Knowledge: By examples. *Analysis*, 27:1–11, 1966.

Murali Ramachandran. Anti-luminosity: Four unsuccessful strategies. *Australasian Journal of Philosophy*, 87:659–673, 2009.

Dana Scott. Advice on modal logic. In Karel Lambert, editor, *Philosophical Problems in Logic*, pages 143–173. Reidel, Dordrecht, 1970.

Glenn Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, 1976.

Timothy Williamson. Inexact knowledge. *Mind*, 101:217–242, 1992.

Timothy Williamson. Does assertibility satisfy the S4 axiom? *Crítica*, 27:3–22, 1995.

Timothy Williamson. Cognitive homelessness. *The Journal of Philosophy*, 93:554–573, 1996.

table. Are there other reasons for maintaining that claim? I'm inclined to think that there aren't, but investigation of the matter lies beyond the scope of this paper.

Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, Oxford, 2000.

Elia Zardini. A model of tolerance. *Studia Logica*, 90:337–368, 2008a.

Elia Zardini. *Living on the Slippery Slope. The Nature, Sources and Logic of Vagueness*. PhD thesis, Department of Logic and Metaphysics, University of St Andrews, 2008b.

Elia Zardini. Towards first-order tolerant logics. In Oleg Prozorov, editor, *Philosophy, Mathematics, Linguistics: Aspects of Interaction*, pages 35–38. Russian Academy of Sciences Press, St Petersburg, 2009.

Elia Zardini. Luminosity and determinacy. *Philosophical Studies*, 2012a. Forthcoming.

Elia Zardini. First-order tolerant logics. *The Review of Symbolic Logic*, 2012b. Forthcoming.