LUMINOSITY AND DISPOSITIONS TO BELIEVE

Iñaki Xavier LARRAURI PERTIERRA

ABSTRACT: Defences of Williamson's Anti-Luminosity Argument (ALA) that employ doxastic propagation principles—i.e., rules by which cases of beliefs and/or dispositions to believe are inferred from other such cases—risk running into sorites. Since these principles are explainable by an ineffective capacity to phenomenally discriminate between two adjacent cases, luminist rejections of the ALA can halt sorites by denying doxastic propagation, thereby reaffirming these discriminative capacities as appropriately effective. One potent method of resisting the luminist involves recharacterizing discriminative capacities in terms of a distinction between beliefs and their underlying dispositions to reinstate the plausibility of doxastic propagation. To this effect, I propose a novel coarse-grained approach favouring the ALA that leverages a modal analysis of the belief/disposition distinction. This motivates a sharp threshold between belief and absent belief that neither succumbs to sorites nor begs the question against the luminist. The upshot of this approach is substantial: by conferring a dialectical advantage to the anti-luminist, the luminist is held to highly problematic positions, regarding the trivialisation of safety and the relationship between beliefs and dispositions, if they deny the coarse-grained approach at many of its components.

KEYWORDS: Williamson, luminosity, margins principle, disposition to believe, sorites

Luminosity is the thesis that if a proposition is true, then one is at least in a position to know that it is true. Famously, Williamson's (2000) *Anti-luminosity Argument* (ALA) denies this. One ramification of this denial is the establishment of an epistemic margins principle, whereby knowledge is factive for the actual knowledge-case and for nearby cases. This is because margins principles directly contradict luminosity in how to evaluate situations in which a case of propositional knowledge neighbours another case of propositional falsity: luminosity allows for such situations while margins principles explicitly forbid it.

However, depending on how margins is motivated, the ALA runs the risk of sorites. This is especially true when the motivation relies on principles of belief propagation—i.e., principles by which cases of belief are inferred from other cases of belief—for then there is a risk of inferring belief within cases wherein it is clearly false—e.g., believing that you are cold while in the centre of the sun. Even so, such propagation principles are intuitively applicable to human knowers given our imperfect capacities to phenomenally discriminate between cases that are similar to each other, hence the relevance of propagation principles and the persistence of beliefs past just the actual case of belief. Therefore, the anti-luminist is compelled to account for propagation principles in ways that are sensitive to when beliefs stop propagating. This is why a plausible luminist rejection of margins can appeal to these propagation principles being either *ad hoc* in their sensitivity to cases of absent belief or simply soritical. This remotivates at least the plausibility of some propositions being luminous. Indeed, some luminists—e.g., Barz (2017)—do consider that our discriminative capacities are effective enough for some propositions that these propositions can be correctly deemed luminous. Such is the case when Berker (2008) commented that luminosity obtains on the grounds that we possess a form of "doxastic privileged access" (2, italics removed) to some facts of the matter that condition particular propositions as luminous.

This partly explains why recent literature on the topic has focused on the plausibility of doxastic models of propagation. For instance, some anti-luminists, like Srinivasan (2015), analyse a distinction between beliefs and their underlying dispositions to show how, given appropriate interpretations of what counts as a disposition to believe, propagation of beliefs and dispositions not only is independently plausible but also does not risk sorites. On the other hand, luminists like Vanrie (2020) have criticised such models by noting how some of their components are not motivated enough and/or not independently well-established.

This paper gives a coarse-grained ALA analysis in terms of a belief/disposition distinction that places the focus on certain modal aspects surrounding actual cases of belief, dispositions to believe, and knowledge. The purpose of this focus is to clarify a sense of belief and its underlying disposition that is not susceptible to sorites while being independently plausible on account of not begging the question against more luminist intuitions about the effectiveness of our discriminative capacities. To accommodate these demands, I clarify a distinction between those cases throughout which beliefs and/or dispositions propagate *because* they are mutually indiscriminable, those within which propagation can be halted *despite* them being so indiscriminable, and those within which propagation is halted *because* they are mutually discriminable. A move to a finer-grained analysis is therefore unnecessary for the anti-luminist if a coarse-grained resolution can be deployed instead.[1]

[1] The difference between a fine- and coarse-grained analysis is that, generally, a fine-grained analysis of the ALA conceives of specific variables as changing in value along a continuum across a series of cases. These variables can be, one, propositions themselves, specifically their content and/or truth values (Williamson 2000, 105ff.; Vogel 2010, 555ff.), or two, the strength our beliefs (Srinivasan 2015, 309, 313-314). A coarse-grained treatment, alternatively, also deals with the

My analysis here is novel on two accounts. First, it introduces two dimensions—one modal and another discriminative—that make better sense of, one, the positionality of cases within the case-series relevant to the ALA, and two, a distinction between when beliefs and when dispositions start falling off along such series. These dimensions are, one, a capacity for non-phenomenal, *dispositional* discrimination between cases, and two, a modal feature of *case neighbourhood*, as opposed to sufficient similarity, between cases. Insofar as intuitions about sufficient similarity and phenomenal discriminability are insufficient for a comprehensive analysis of the ALA in terms of doxastic propagation, it would help to introduce more dimensions to better appreciate the sense in which doxastic propagation contributes to a fuller ALA analysis. I also argue that this way of analysing the ALA via these two dimensions grants unique conceptual resources for resisting the luminist's charge that knowledge does not require epistemic margins principles because cases of knowledge are undergirded by 'doxastic privileged access' to known propositions. Second, my analysis affords significant dialectical advantage to the anti-luminist, since denying it forces one to espouse quite problematic claims concerning safety and/or the relation between beliefs and their underlying dispositions. I accomplish this on the back of the propagation rules that establish margins because, by denying these rules, one must thereby accept that beliefs can be trivially safe due to dispositions either equating or having nothing to do with beliefs.

I begin by briefly outlining an ALA in terms of phenomenal discrimination to motivate the concepts of phenomenal indiscriminability and case neighbourhood (Section I) before expanding upon how the margins principle the ALA motivates can be further substantiated by these concepts (Sections II&III). I then introduce the belief/disposition distinction (Section IV), analyse how through it one can establish a margins principle without risking sorites (Section V), and further flesh the analysis out by applying the notion of dispositional indiscriminability between cases (Section VI). I finally explore how all this dialectically serves the anti-luminist against the luminist (Section VII) before evaluating how a possible luminist rejoinder ultimately fails to settle the debate in the luminist's favour (Section VIII). I conclude in Section IX.

I

The ALA that I outline is based closely on Williamson's (2000) own version:

---

changing in value of either propositional truth/content and/ or belief strength across a case-series, although the main difference is that this change does not happen gradually—i.e., there is a sharp point within a series of cases where the state of affairs suddenly changes from one having a belief to one not having one, or from a proposition being true to it being false.

> We set up a temporal series of cases, $\alpha_{0\text{-}n}$, between t_0 and t_n, composed of mutually sufficiently similar case-pairs—e.g., (α_i, α_{i+1}) represents a sufficiently similar case-pair because it is composed of mutually sufficiently similar cases. The content of the proposition, C, changes gradually throughout this series such that C obtains in α_0 but not in α_n.[2] Furthermore, C is such that, if C obtains in α, then, in α, S is in a position to know that C (i.e., K_{pos}C)—call this the luminosity principle, L, for any luminous proposition, C. S's epistemic capacity is also such that, $\forall(\alpha_i : 0 \leq i \leq n)$, if, for S, K_{pos}C obtains in α_i, then in α_{i+1} C obtains—call this the margins principle, R, as applied to luminous propositions. If C obtains in α_0, then, given both L and R, C must also obtain in α_n.[3] This of course contradicts the assumption that C does not obtain in α_n, so if the first two assumptions must stay in place given the set-up of the scenario, then one must reject either L or R; but R is merely a principle of our inexact knowledge—it follows naturally from the idea that there are cases between which the truth value of C is phenomenally indiscriminable for us, thereby making it implausible for one to be able to possibly know something in one case when, say, in an immediately adjacent case they could easily be claiming that something is true when in fact it is not—so L must go, or so the anti-luminosity argument says.

Broadly speaking, phenomenal indiscriminability licenses inferences from belief-cases to other relevant belief-cases, where what makes a belief-case relevant depends on how indiscriminability is conceived. For instance, Vogel (2010) describes a doxastic principle called indiscriminability-3, whereby, "[g]iven that cases *α*, [*α*'] are indiscriminable-3 to one, if in *α* one believes that C obtains, there is some case [*α*''] relevantly just like [*α*'], in which one believes that C obtains." (562) There are of course other ways of cashing out the terms of indiscriminability, (564; see also, Williamson 2008) but for our purposes here, U being phenomenally indiscriminable from V entails that one is liable in U to believe that C if and only if one is liable in V to also believe that C. One's 'liability' to believe can be interpreted in terms of

---

[2] C's propositional content is specifically the feature of the case to which C refers and is what matters for the truth value of C in that case. We may even say that C's content in α_0 verifies C in α_0. For example, if C is, 'it is cold', then its content would be either the specific temperature or some cold phenomenon to which 'it is cold' refers. Alternatively, if C is, 'I feel cold', then its content would be one's specific cold feelings. Now, ¬C would be false in α_0, but ¬C would *still have the same propositional content as C*, for content is properly a feature of the case, not the proposition itself, meaning that the content both verifies C and falsifies ¬C. We make this distinction between proposition, content, and case simply to illustrate the relevant particularities of the ALA.

[3] The focus is on luminous propositions, which are differentiated from beliefs themselves. It may be the case that beliefs, or judgments more specifically, are luminous and hence are not associable with a margins principle—Soteriou (2013) and Jenkins (2021) claim as much—but they are not the primary topic of this paper. However, as implied in Sections VII-VIII, what the luminist may need to admit to have beliefs and judgments be luminous may be problematic.

indiscriminable-3, but 'liability' here is taken more generally to be based on a type of content relation between cases that matters to how well one can discriminate one from the other, at least in terms of the truth-value of the proposition in question—hence why the ALA is partly conceived in terms of propositional content. This relation is fleshed out in the succeeding paragraphs.

With that being said, we can generalise some of the ALA's components. L as stated is explicitly talking about luminosity as a principle concerning one being in a *position to know* (Williamson 2000, 95); however, for our purposes, if we assume that between cases α_0 and α_n S "thoroughly considers" whether C obtains "throughout the process, we can ignore the difference between knowing and being in a position to know." (Wong 2008, 537; see also, Williamson 2000, 95) This just means that, supposing a belief requirement for knowledge is in place, throughout the progression of cases, $\alpha_{0\text{-}n}$, it is true that $K_{pos}C \rightarrow BC$. Thus, if the difference between S knowing that C (i.e., KC) and $K_{pos}C$ is simply S believing that C (i.e., BC) on the basis of their epistemic position regarding C, then we have $(K_{pos}C \wedge BC) \rightarrow KC$, and therefore $(K_{pos}C \rightarrow KC)$ applies throughout $\alpha_{0\text{-}n}$.[4] Here, L can be restated as,

$$L^* \Leftrightarrow (C \rightarrow KC).$$

R as stated currently just refers to knowledge in one case being factive over this and a temporally adjacent *future* case. (Williamson 2000, 97-98; Brueckner & Fiocco 2002, 286) By keeping this temporal feature of R,we can otherwise modify it to specify how the relevant epistemic and propositional truths are based. If we write 'C obtains in α_i with propositional content *a*' as '$\alpha\text{-}C_i^{a}$', then, $\forall[(a, b, \alpha_i) : (|a{-}b| \leq m), (0 \leq i \leq n)]$,

$$R^* \Leftrightarrow [\alpha\text{-}(KC)_i^{a} \rightarrow \alpha\text{-}C_{i+1}^{b}],$$

where *m* represents a margin within which the truth value for C is *phenomenally* indiscriminable. (Williamson 2000, 115; Dokic & Égré 2009, 3) '$\alpha\text{-}(KC)_i^{a}$' therefore expresses an instance of 'knowledge that C' in α_i alongside propositional content, *a*. Here, *a* is not properly considered as content that directly matters in α_i for KC, because *a* properly is content that matters in α_i for C. However, this does not mean that *a* and KC exhibit no relation, for while *a* does not single out any specific mode of justification supporting KC, it can be involved in however KC is justified—e.g., if

[4] Ultimately, this difference between knowing and being in a position to know is immaterial when it comes to counterexamples against luminosity in terms of phenomenal indiscrimin*ability*, which Williamson (2021, 108) notes in relation to counterexamples against the KK-thesis, although we should be able to generalise this against luminosity claims more broadly.

C is, 'it is cold', then *a*, representing the temperature or cold phenomenon verifying C, surely matters for knowing whether it is cold or not.[5]

Importantly, not only do *a* and *b* represent different instances of propositional content that still verify C in both α_i and α_{i+1}—e.g., *b* does not express content that falsifies C in α_{i+1}—but that it is possible for these different instances to be expressive of the same content, which occurs when $|a - b| = 0$. The relation, *m*, expresses the notion that the instances of propositional content for C in both α_i and α_{i+1} are such that, if *b* instead verifies ¬C while *a* verifies C, the value of *m* would prohibit S from discerning whether C or ¬C obtains in either α_i or α_{i+1}. What is significant here, though, is that *m* represents one's margins for *phenomenal* indiscriminability, so not all facets of propositional content in *a* and *b* matter for *m*, just the features of said content that play a role in how C is phenomenally expressed in α_i and α_{i+1}, respectively, such that one is liable to believe the same thing in both cases. Consequently, *a* is not equivalent to its phenomenal content, so we can also talk about *non*-phenomenal modes of indiscriminability whereby what makes one liable to believe C in two indiscriminable cases has to do with features of its content that play a role in how C is non-phenomenally expressed in those cases. Regardless, this only becomes relevant once we begin analysing the significance of *non*-phenomenal indiscriminability for doxastic dispositions in Section VI. For now, we stick to discussing phenomenal indiscriminability between cases.[6]

Suffice it to say, talk about indiscriminability here is kept rather general, only involving the broad relation of *m*, to not preclude from the outset entailment rules capturing different ways in which indiscriminability could play out doxastically. As we will see, some of these rules become integral to the establishment of margins principles like R* and have analogous precedent in the literature, so keeping the discussion of indiscriminability sufficiently broad suits our purposes here. Furthermore, the discussion is also broad enough to not presume a specific relation between a case's content and S's (i.e., agent's) discriminative capacity. This is to not beg the question against either internalists or externalists about content. This means that *m* should remain useful in our discussion whether it is agent- or case-indexed—i.e., whether *m* shifts in value in accordance with S, the case in question, or a combination of both.[7]

---

[5] This is not to say that *a* is what *only* matters for KC; other factors may come into play as well, such as the temperature/cold phenomenon in nearby cases. Moreover, R* can be further generalised to include non-luminous conditions, but R* is sufficient for our present purposes.

[6] Talk of indiscriminability from now on will drop the 'phenomenally' moniker but will still, unless otherwise specified, mean phenomenal indiscriminability.

[7] A content internalist interpretation of the relation between *m* and one's discriminative capacity

What notion the ALA speaks to is that knowledge is not merely factive, but also "factive with respect to neighbouring cases" that fall within the margin, *m*, in question. (Dokic & Égré 2009, 2) It is the notion that, in Brueckner and Fiocco's (2002) interpretation of Williamson,

> if S passes from correctly believing C at t_i to mistakenly believing C at t_{i+1}, then this shows that S's confidence regarding C at t_i (which hardly differs from his confidence at t_{i+1}) *was not reliably based.* Since S's putative knowledge rests entirely on this basis, then because it is incapable of supporting knowledge at t_{i+1}, it must have been so at t_i. (287)

This forms the motivational basis for margins principles like R* expressed in terms of doxastic reliability within neighbouring cases. These neighbouring cases are currently expressed as *actual* temporally adjacent cases, consisting of the actualised temporal series of cases, $\alpha_{0\text{-}n}$. In Section III we explore how a similar margins principle can accommodate a *possible* temporal series of cases, $\beta_{0\text{-}n}$, such that cases that neighbour α_i can either be actual or possible cases temporally adjacent to it. For now, it suffices to say that case-neighbourhood, the property inhering in cases that neighbour each other, and temporal adjacence are used interchangeably. Generally, then, the case-neighbourhood relation is as follows: if 'A neighbours B' is written as 'B $\Leftrightarrow_N$ A', then, considering any general case, δ, we have, $\forall[(\delta_i) : (0 \leq i \leq n)]$, $\delta_i \Leftrightarrow_N$

---

entails that *m* is a matter of how *a* and *b* appear such as to be accessible phenomenologically by S. Alternatively, a content externalist reading entails that *m* is a matter of how the phenomenal appearances of *a* and *b* obtains within the cases in question. In other words, when two cases are phenomenally indiscriminable regarding the contents in question, this is a function either of who S is (internalist) or of the cases themselves (externalist), or perhaps a combination of both. It is more intuitive to talk about discriminative capacities directed at phenomenal content in internalist terms, given its connection with phenomenal appearances. Interpreting such talk in externalist terms is less intuitive albeit still plausible—e.g., S's incapacity to phenomenally discriminate could be causally connected to agent-independent features of cases that bear on how *a* and *b* appear phenomenologically to S—although this may be more viable when associated, as expounded upon in Note 34, with *non*-phenomenal content and *dispositional* indiscriminability. This externalist picture is reminiscent of Sosa's (2015) distinction between complete and incomplete epistemic competences, wherein a complete competence is so by virtue of being sensitive to a range of salient situational/environmental features involving the epistemic agent. An analogous internalism would have it that phenomenal content itself is directly linked to one's incapacity to phenomenally discriminate without any causal involvement with separate case features. In other words, either our incapacity to phenomenally discriminate is grounded on phenomenal indiscriminability as a function of S-independent case content (externalism) or phenomenal indiscriminability about case contents is S-dependent by being grounded on S's relative incapacity to phenomenally discriminate (internalism).

$\delta_{i\pm1}$. Later in this Section we broaden the case-neighbourhood relation to better conceive of the connection between the ALA and non-neighbouring cases.

Now, some have found issue with R*, or margins principles like it generally,[8] and it is easy to see why. For instance, let us consider a safety requirement for knowledge where, to stick with the language of R*, knowledge necessitates the absence of false beliefs in neighbouring and indiscriminable cases to and from the actual case in which knowledge obtains, respectively. In other words, this safety requirement, call it **A***, has it that,

$$\mathbf{A^* \Leftrightarrow [\alpha\text{-}(KC)_i{}^a \rightarrow (\alpha\text{-}(\neg BC)_{i+1}{}^b \vee \alpha\text{-}C_{i+1}{}^b)].}$$

A* can also be motivated on Williamson's (2000) own conception of safety, since he conceives of doxastic reliability in terms of one's "confidence" in nearby cases attaining "a similar basis" to "one's almost equal confidence" in the actual present case. (97; see also, 101, 149) Here, $\alpha\text{-}(BC)_{i+1}{}^b$ represents a 'belief that C' in α_{i+1} alongside C's propositional content, *b*, and *b* is very similar to *a*, which is the content relevant to one's 'belief that C' in α_i.[9]

However, A* being accepted does not give R* for free. The crux is that if we, like Williamson (1996; 2005), regard knowledge itself as non-luminous—e.g., it is possible for $\alpha\text{-}(KC)_i{}^a \wedge \alpha\text{-}(\neg KKC)_i{}^a$ to obtain—then requiring both A* *and* R* for knowledge would force us to guarantee $\alpha\text{-}(KC)_i{}^a$, $\alpha\text{-}C_{i+1}{}^b$, and the possibility of these three conditions all at once: $\alpha\text{-}(BC)_{i+1}{}^b$, $\alpha\text{-}(BC)_{i+2}{}^c$, and $\alpha\text{-}(\neg C)_{i+2}{}^c$, where $0 < |b - c| \leq m$. This substantiates Williamson's (2000) claim that "[o]ne can be reliable without being reliably reliable" regarding knowledge, (125) because one's reliability and lack of reliable reliability would then be due to α_i, α_{i+1}, and α_{i+2} representing, respectively, a series of knowledgeable, true but unreliable, and false beliefs by virtue of $|a - b|$ and $|b - c|$ both being within one's margins, *m*.[10] To connect A* to R*, what is needed

[8] Dorst (2019) comes to mind when they argue that accepting margins for knowledge forces one to make infelicitous assertions (1250-1251). McHugh (2010, 252-255) and Das and Salow (2018, 20n10) consider particular modes of knowledge as requiring no margins principle. Finally, on a more margins-considerate angle, while Immerman (2018, 413-414; 2020, 3369-3371) argues for the inapplicability of margins principles in situations where the right doxastic bases are at play, Bonnay and Égré (2008, §5.3.1) and Ramachandran (2012, 122) all contend that, while knowledge may have margins, it would be highly irrational to expect anyone to *know* what those margins are.

[9] Since *a* properly is a feature of α_i, and not of the instance of knowledge or the proposition itself (i.e., *a* can obtain regardless of C or ¬C being true), $\alpha\text{-}(BC)_{i+1}{}^b$ represents an instance of BC alongside propositional content, *b*, that is a feature of the case, α_{i+1}, not of the belief itself.

[10] Note that $|b - c| \neq 0$ because C's truth-value difference ought to be a sufficient condition for its difference in propositional content if the latter matters to the truth value of C. Nevertheless, this should not affect $|b - c| \leq m$ given that content differences in neighbouring, *indiscriminable* cases must still be within one's margins for *phenomenal* indiscrimination, which is consistent with those

then is something like a principle of belief propagation along a case-series, because, as Jenkins (2021) notes, "the claim that knowledge is undermined by a nearby possibility of falsehood" would need to be substantiated by the principle that, "had that possibility [of falsehood] obtained, one would have been in imperceptibly different circumstances from one's point of view to those in which one is in [a belief-preserving manner]" (S1558)—e.g., supporting R* by both A* *and* something like the possibility of $\alpha\text{-}(\neg C)_{i+2}{}^{c}$ being irrelevant for the move from $\alpha\text{-}(BC)_{i+1}{}^{b}$ to $\alpha\text{-}(BC)_{i+2}{}^{c}$ (Srinivasan 2015, 295, 301). What this looks like is illustrated succinctly by Berker (2008), who contends that,

> Williamson's [ALA] only succeeds if he assumes that we do not have a kind of *doxastic* privileged access (as we might put it) to the facts in question, for his argument presupposes that there does not exist a certain sort of constitutive connection between the obtaining of the given facts and our *beliefs* about the obtaining of those facts. (2)

This idea that a lack of 'doxastic privileged access' to facts like content differences between C and ¬C supports R* implies that there is something about the cases relevant to safety (A*) that matters to the indiscriminability of C's truth value between those cases in a way where they also become cases relevant to margins (R*).

However, before doing so we must relate the terms used by both A* and R* to guard against any potential dimension in which both are mutually irrelevant. R* deals with the neighbourhood of indiscriminable cases, while A* does not have to be couched in these terms. Indeed, formulations of safety commonly have two relational dimensions: a modal dimension of sufficient similarity that fixes the nearness relation between different cases, and a discrimination dimension whereby safety does not apply to mutually discriminable cases. In other words, it is usual for a belief's safety to require this belief to "not be false in any similar case that one cannot discriminate from [the actual present case]" (Berker 2008, 3). Similarity and discriminability are related since cases may be mutually indiscriminable if they are sufficiently similar to each other, especially if they are similar in phenomenal content. However, the most obvious difference is that cases can be mutually sufficiently similar, or indiscriminable, without ever neighbouring each other, and *vice versa*.[11]

---

differences effecting propositional truth-value differences in indiscriminable ways.

[11] There are also examples in the literature of conceiving of safety in terms of something more positional, like what case-neighbourhood represents. For instance, Smithies (2012) notes that one "can know that C obtains, even if C does not obtain in every *close* case, so long as there is no *close* case in which [one] falsely believe[s] that C obtains." (725-726. Emphasis added)

To resolve these differences between similarity, indiscriminability, and case neighbourhood, let us construct viable relations between them. However, one cannot simply have neighbouring and mutually indiscriminable cases be also mutually sufficiently similar, for this risks, at least in the context of the ALA, either wedding similarity merely to discriminability and/or case neighbourhood or denying temporally adjacent cases ever being mutually *dis*similar. Both are problematic in the broadest sense of being *ad hoc*. Indeed, having case neighbourhood be reduced to some combination of sufficient similarity and indiscriminability would not address the above concerns. For instance, we know that α_i and α_{i+1} are neighbouring case-pairs because they are temporally adjacent to each other, but this information alone does not denote extra knowledge of their sufficient similarity and/or indiscriminability: someone not being able to discriminate between C in α_{i+1} and ¬C in α_{i+2} is a function of *m*, which is a content relation that matters for how mutually similar these cases are, and not of them being mutually neighbouring—after all, (α_i, α_{i+1}) being a sufficiently similar case-pair is not necessary as it had to be assumed from the outset. If the difference in content between α_{i+1} and α_{i+2} is greater than *m*, then them being a neighbouring case-pair would be irrelevant to them being mutually discriminable.

Still, case neighbourhood and similarity are not completely independent: α_i neighbours α_{i+1} but not α_{i+2}, and we would expect (α_i, α_{i+1}) to be a more similar case-pair than (α_i, α_{i+2}). Additionally, the property of these cases being 'neighbouring' presumably captures what is significant to the ALA: intuitively, cases that are sufficiently similar enough to each other to also be mutually indiscriminable instantiate primarily as those adjacent, *neighbouring* case-pairs that constitute the actual case-series of the ALA. More generally though, to not beg the question against the luminist—i.e., to allow for mutually discriminable and neighbouring case-pairs—the more direct connection would have to be between sufficient similarity and case neighbourhood. As such, although neither indiscriminability nor sufficient similarity are necessary conditions for case neighbourhood, since neighbouring cases may be sufficiently *dis*similar enough to where they become mutually discriminable regarding some content difference that is larger than *m*, it may be that as cases become increasingly dissimilar the possibility for them to instantiate as neighbouring pairs diminishes.

This entails the possibility for some content similarity threshold past which cases enter into a kind of neighbouring relation. Now, a case-pair meeting this threshold ought not guarantee case neighbourhood, for that assumes too much. However, it does not seem *prima facie* impossible that, for any sufficiently similar case-pair that is non-neighbouring, there is some possible case-pair, attaining the same content relation as the first pair, that is neighbouring. To this end, I introduce a relation called,

> **NeighSim**: a case-pair's sufficient similarity meets some threshold, *s*, only if it either is a neighbouring case-pair or involves a content difference whose value is identical to a possible neighbouring case-pair.

To help distinguish, a neighbouring case-pair possesses the relation, $\Leftrightarrow_{N}$, while a case-pair that either is neighbouring or involves content identical to a possible neighbouring case-pair possesses the relation, $\Leftrightarrow_{N\backslash PN}$. Cases that either neighbour each other or share content with those that could neighbour each other are called N\PN-related cases.

NeighSim correctly entails that sufficient similarity and case neighbourhood are separate relations, in that any non-neighbouring case-pair must be insufficiently similar only if it is impossible for there to be an alternative (content difference)-identical case-pair that is possibly neighbouring—but this impossibility seems deeply unintuitive on account of it being an incredibly strong claim requiring there to be no extent of A) case similarity and/or B) indiscriminable content difference that any non-neighbouring case-pair could exactly share with even a possibly neighbouring one. Additionally, the fact of this impossibility being problematic also affords the NeighSim supporter a response to the sceptic: the possibility of a neighbouring indiscriminable case-pair consisting of a sceptical and the actual case, which would defeat our knowledge in the actual case, is undermined because the sceptical case would either have to neighbour the actual one by being temporally adjacent to it, which is absurd if the actual case is not already a sceptical case, or it would have to be *both* sufficiently similar to the actual case *and* express a content relation to it that is impossible to find in another content-identical, possibly neighbouring case-pair, which, as implied above, is deeply problematic.[12]

---

[12] We can concede, though, that sceptical scenarios may be N\PN-related to α_i without neighbouring nor being sufficiently similar to it. This can deny us knowledge in α_i via a safety requirement for knowledge wherein the condition of neighbouring, indiscriminable cases is replaced with one of N\PN-related, indiscriminable cases. This is a possibility only because the *a*-content of $\alpha\text{-}C_i^{a}$ does not exhaust specification of α_i, so a sceptical case can be N\PN-related to α_i if it shares its content with a non-sceptical case that either neighbours or possibly neighbours α_i. The observations by Weber and Omori (2019, 990ff.) allow for something analogous regarding cases that both are mutually insufficiently similar and express an asymmetric accessibility relation, for they note a general unsettledness in determining where an actual case sits in relation to impossible ones within attempted representations of modal space. However, given that case content is not just phenomenal, it has yet to be determined whether there can even be a sceptical case that is content-identical to a non-sceptical one since, intuitively, a sceptical case's content should matter *to it being a sceptical case*.

We can now, by this relation of sufficient similarity with case neighbourhood through NeighSim, consolidate the terms for R* and A*. By having both be sensitive to mutually indiscriminable and N\ PN-related cases, we not only make both rules more in line with each other in their applicability to the ALA, but also reasonably extend how safety is commonly deployed in the literature. Additionally, this change more clearly addresses a plausible doxastic rationale underlining both rules, in that our attenuated capacity for discrimination should be insensitive not only to our *actual* neighbouring cases but also to those that *could* neighbour our own present case. Margins and safety principles stem from our inability to discriminate between nearby cases since, one, by this inability we would be liable to have the same belief in those cases, and two, a guarantee of absent false belief for knowledge must be grounded elsewhere if auspicious discriminative capacities are not forthcoming. This guarantee is therefore, to be more sensitive to N\PN-related cases, fulfilled in part by an auspicious modal landscape instead, where cases of false belief are posited as far off possibilities to not undermine actual instances of knowledge. Nevertheless, it would be very strange if this landscape consisted solely of either sufficiently similar and/or neighbouring cases. This is because, one, case neighbourhood, more than sufficient similarity, is a more direct measure of modal closeness/nearness given its already built-in positionality metric, 'i', and two, case neighbourhood alone cannot fully accommodate the ramifications of having an attenuated discriminative capacity since one can surely be unable to discriminate one case from not only a neighbouring one but also one that could *possibly* be neighbouring.

Consequently, $\forall(a, b, \alpha, \beta, \delta, \mathrm{i}, \mathrm{x}) : (\{\alpha, \beta\} \in \{\delta\}), (\alpha_{\mathrm{i}} \Leftrightarrow_{\mathrm{N\backslash PN}} \delta_{\mathrm{x}}), (|a - b| \leq m), (0 \leq \mathrm{i} \leq \mathrm{n})$,

$$\mathbf{R+} \Leftrightarrow [\alpha\text{-}(\mathrm{KC})_{\mathrm{i}}{}^{a} \rightarrow \delta\text{-}\mathrm{C}_{\mathrm{x}}{}^{b}]$$

$$\mathbf{A+} \Leftrightarrow [\alpha\text{-}(\mathrm{KC})_{\mathrm{i}}{}^{a} \rightarrow (\delta\text{-}(\neg\mathrm{BC})_{\mathrm{x}}{}^{b} \vee \delta\text{-}\mathrm{C}_{\mathrm{x}}{}^{b})]$$

'x' is meant to label any general (δ) case, actual (α) or possible (β), that is N\PN-related to α_i. There is no absolute restriction on the numerical value for 'x', but generally, $\forall(\mathrm{i, x, y, z}), (0 \leq \mathrm{i} \leq \mathrm{n}), (\mathrm{i\text{-}y} \leq \mathrm{x} \leq \mathrm{i+z}), (1 \leq \mathrm{y} \leq \mathrm{i+n+1}), (1 \leq \mathrm{z} \leq \mathrm{n\text{-}i+1})$. Now, R+ is already much stronger of a claim than what the anti-luminist needs to establish if all they need is a margins principle that accommodates the ALA as presently set up. Also, α_{i+1} is N\PN-related to α_i anyways, R+ begs too much compared to R* against the luminist concerning the extent to which our discriminative capacities are attenuated, and the anti-luminist can already respond to luminist critiques of R* without having to respond to those of R+. As is shown below, the anti-luminist can adopt rather modest principles of belief propagation that already motivate R* without committing them to the problem of defending the ALA by requiring belief of the same proposition within more mutually indiscriminable cases than needed. Resultantly, I will only be discussing R* moving forward. A+, on the other hand, can be easily targeted by the luminist to undermine R* as long as they can guarantee a result like $(\alpha\text{-}(KC)_{i+1}{}^{b} \wedge \alpha\text{-}(\neg C)_{i+2}{}^{c})$ by appealing to some other safety principle more motivated than A+. However, as is argued in Section VIII, this ability of the luminist to appeal to alternative safety principles can also be adopted by the anti-luminist to reestablish R*. To set up how this manoeuvre is warranted in the later Sections, I will continue considering A+ moving forward.

This is not to say that A+ cannot be independently motivated, as it respects intuitions about safety dealing with both possible and actual cases. For instance, Srinivasan (2015) notes that, putatively,

> [t]o pass the safety test for knowledge, it is insufficient that one, as a matter of chance, lack untrue belief in all actual similar cases. One must also lack untrue belief in *possible* similar cases. This means that it is much easier for S to fail to know that she is cold than some luminists seem to think. (302)

Adapting this thought into A+ is rather easy. We adopt the usual terminology whereby β_i is a possible case for α_i and β_{i+1} is a possible case for α_{i+1}, and where β_{i+1} (or β_i) is not necessarily possible for α_i (or α_{i+1}). The relation of 'possible for' is distinct from that of indiscriminability, so it is possible for β_{i+1} to both be indiscriminable from α_i and not be a possible case for it.[13] A possible case must be N\PN-related to α_i if it is to also be the general case, δ_x. Now, β_{i+1} does not neighbour α_i in the sense of *actually* being temporally adjacent to it, but it can certainly be N\PN-related case to it. That is to say, a (β_{i+1})-case is a (δ_x)-case if it *could* actually neighbour α_i, if what specifies β_{i+1} could also specify some actual case temporally adjacent to α_i, where the non-neighbouring (α_i, β_{i+1}) case-pair would therefore attain a content relation identical to some possible neighbouring case-pair—i.e., where $(\alpha_i \Leftrightarrow_{N\backslash PN} \beta_{i+1})$ would obtain. Note also that all this can establish $(\alpha_i \Leftrightarrow_{N\backslash PN} \beta_i)$ as well. We thus have a rule that will apply going forward: $\forall[(\alpha, \beta, \delta) : (\{\alpha_{i+1}, \beta_i, \beta_{i+1}\} \Leftrightarrow_{N\backslash PN} \alpha_i)], \{\alpha_{i+1}, \beta_i, \beta_{i+1}\} \in \{\delta_x\}$.

II

The task then is to convincingly motivate R*. We cannot abide by A+ alone for this task not just for the fact that R* cannot be straightforwardly derived from A+ but

[13] While this would be problematic in the sense of '$\alpha_i \Leftrightarrow_N \beta_{i+1}$', for it would not make sense to call a neighbouring case to α_i impossible for it, it *could* make sense if both cases were N\PN-related instead—i.e., $\alpha_i \Leftrightarrow_{N\backslash PN} \beta_{i+1}$. However, see Note 12. Notwithstanding, A+ requires *possible* true belief, not an *impossible* one, so only those β_{i+1} cases possible for α_i are considered when we make use of A+.

also because both paint different pictures of knowledge. Williamson (2000) couches indiscriminability as a function in part of a truth insensitivity, or, a doxastic insensitivity to a concept's own truth conditions, whereby there is "limited discrimination in the [corresponding] belief-forming process" (127; see also, 103-104) concerning these differences in the instantiation of such truth conditions. (see also, Williamson 2021, 108) If we accept this, then, by relying only on A+, one permits the coinciding of knowledge and N\PN-related indiscriminable (¬C)-verifying cases as long as one does not in such cases believe that C. However, these are precisely the scenarios disallowed by R*, so espousing both A+ and R* entails accepting that knowledge can obtain only on one of two conditions: our attenuated discriminative capacities are either, one, capacious enough that no (¬C)-verifying case is indiscriminable from α_i, or two, incapacious enough that knowledge needs to be supported by an auspicious enough modal landscape wherein, if such cases are indiscriminable from α_i, then they could not even possibly neighbour it. The first condition concedes too much to the luminist, so a viable anti-luminist motivation for R* must consider the second one.

One way to cash out this second condition would be something like the doxastic propagation principle,

$$B+ \Leftrightarrow (\alpha\text{-}(BC)_i{}^a \rightarrow \delta\text{-}(BC)_x{}^b).$$

However this, as Berker (2008, 7, 17) correctly observes for propagation principles like B+, leads to a sorites when repeatedly iterated upon. He also considers that, "as it is incontestable that BC does not obtain in case α_n," Williamson cannot appeal to epistemic safety conditions to derive "the conditional 'KC obtains in α_{n-1} → C obtains in α_n', which is one of the instances of [R*]" (8), as long as we take the conditional to also imply the proper propositional content relation, $|a - b| \leq m$, between cases α_{n-1} and α_n—i.e., the conditional would become, $\alpha\text{-}(KC)_{n-1}{}^a \rightarrow \alpha\text{-}C_n{}^b$, which Williamson cannot derive, due to $\alpha\text{-}(\neg BC)_n{}^b$ obtaining, by appealing to epistemic safety conditions like A+. This also potentially speaks against B+ and principles like it, considering that $[\alpha\text{-}(BC)_{n-1}{}^a \wedge \alpha\text{-}(\neg BC)_n{}^b]$ is one way to falsify B+.

Now, before moving on, I will say that Williamson can accept Berker's argument while still having an out if he appeals to the contrapositive of B+. If Berker simply means to highlight $\alpha\text{-}(\neg BC)_n{}^b$ but *not* $\alpha\text{-}(\neg BC)_{n-1}{}^a$, then this would be consistent with Williamson's appeal to safety conditions despite Berker supposing otherwise, given that, along with $\alpha\text{-}(\neg KC)_{n-1}{}^a$ obtaining, this would also be consistent with $\alpha\text{-}(\neg BC)_n{}^b \rightarrow \alpha\text{-}(\neg BC)_{n-1}{}^a$ being the case. Otherwise, Berker specifying instead that $\alpha\text{-}(\neg BC)_n{}^b \wedge \alpha\text{-}(BC)_{n-1}{}^a$ is the case would simply be begging the question against Williamson's understanding of phenomenal indiscriminability. In short, Berker is saying nothing for which R*'s establishment above cannot already account, such as, if 'it is incontestable that BC does not obtain in case α_n', then by virtue of B+'s contrapositive and a belief requirement for knowledge, it ought to be incontestable that KC does not obtain in case α_{n-1} either. Therefore, the force of the ALA is not necessarily undermined.

The salient point here, however, and one that I believe Berker would need to imply for his above statements to resist Williamson's ALA and not be consistent with it, is that one way to undermine the argument's force is to have R* and B+ cease to apply at some time in the progression between α_0 and α_n. If Berker allows for that time to be between α_{n-1} and α_n such that $[\alpha\text{-}(\neg BC)_n \wedge \alpha\text{-}(BC)_{n-1}]$ obtains, then that should also lead to$[\alpha\text{-}(\neg C)_n \wedge \alpha\text{-}(KC)_{n-1}]$ being true as well—otherwise the ALA could still go through to motivate $\alpha\text{-}(\neg KC)_{n-1}$ instead given R* and an anti-luminist account for $[\alpha\text{-}(\neg BC)_n \wedge \alpha\text{-}(BC)_{n-1}]$. Note that none of this falsifies A+, so we resultantly have a safety-friendly epistemic scenario that contravenes B+'s motivation for R*, and thus for the ALA.[14]

One viable way to safeguard R* would be by permitting $[\alpha\text{-}(\neg BC)_n \wedge \alpha\text{-}(BC)_{n-1}]$ while precluding $[\alpha\text{-}(\neg C)_n \wedge \alpha\text{-}(KC)_{n-1}]$. In other words, for the case series,α_i to α_{i+2}, the needed results would be as follows:$\alpha\text{-}(\neg C)_{i+2}{}^c$, $\alpha\text{-}(\neg BC)_{i+2}{}^c$, $\alpha\text{-}(BC)_{i+1}{}^b$, and $\alpha\text{-}(KC)_i{}^a$, where $|a - b| \leq m$, $|b - c| \leq m$, $|b - c| \neq 0$, and $|a - c| \neq 0$. This is because $\alpha\text{-}(BC)_{i+1}{}^b$ would be true *yet not knowledgeable* belief.[15] Crucially, since B+ is rejected, R* is also consistent with $(\alpha\text{-}(KC)_i{}^a \wedge \alpha\text{-}(\neg BC)_{i+1}{}^b)$, which is more fully explored at the end of Section IV. Nonetheless, for the sake of streamlining the argument, unless made otherwise explicit, we will assume that $\alpha\text{-}(BC)_{i+1}{}^b$ obtains. Consequently, if R* is to not succumb to sorites, then we will need a principle that not only connects $\alpha\text{-}(BC)_i{}^a$ and $\alpha\text{-}(BC)_{i+1}{}^b$, (Srinivasan 2015, 295, 301) but also rationalises $\alpha\text{-}(\neg BC)_{i+2}{}^c$ and $\alpha\text{-}(BC)_{i+1}{}^b$ so that the anti-luminist does not merely beg the question against themselves.[16] As such, what this new principle would need to explain is the

---

[14] Having R* and B+ falsified in this way does not require specification of propositional content relations between $\alpha\text{-}(\neg BC)_n$ and $\alpha\text{-}(BC)_{n-1}$ nor between $\alpha\text{-}(\neg C)_n$ and $\alpha\text{-}(KC)_{n-1}$ other than them being within *m*-bounds.

[15] This case series also concords with Williamson's (2000) own way of constructing a case-series through a principle of conceptual sharpening: "pick *j* and *k* such that $0 \leq j \leq k \leq n$; for each *i*, evaluate 'One feels cold' as true in α_i if and only if $i \leq k$, and otherwise as false; evaluate 'One knows that one feels cold' as true in α_i if and only if $i \leq j$, and otherwise as false" (105). There is an entire literature on conceptual sharpening—e.g., Wong (2008) and Vogel (2010)—that seeks to undermine Williamson's attempt at establishing margins by employing such sharpening. I cannot address it directly here but suffice it to say that the motivations appealed to here to establish R* are at least different to those examined by others in their analyses of Williamson's sharpening of relevant terms like 'feels cold' and 'knowledge'.

[16] $\alpha\text{-}(BC)_{i+1}{}^b$ is significant here when taking its more generalised instance of $\delta\text{-}(BC)_x{}^b$, for (A+ ∧ B+)

distinguishing factor between N\PN-related cases that are indiscriminable—e.g., between α_i and α_{i+1}—and those whose indiscriminability still permits content differences by which propositional and doxastic truth can also differ between the cases—e.g., between α_{i+1} and α_{i+2}.

III

One means of accomplishing this is through a modalized version of B+, called B_M+, where α-$(BC)_i{}^a$ entails the existence of a type of general case, $\delta_{x+1}{}^+$, N\PN-related to α_{i+1} and whose propositional content is identical to the content of α_{i+1} and in which BC is true. In short, $\exists(\delta_x{}^+) : (\{\delta_x{}^+\} \subset \{\delta_x\})$,

$$\mathbf{B_M+ \Leftrightarrow (\alpha\text{-}(BC)_i{}^a \rightarrow [\delta^+\text{-}(BC)_{x+1}{}^b \wedge (\alpha\text{-}C_{i+1}{}^b \equiv \delta^+\text{-}C_{x+1}{}^b)])}.$$

$(\alpha\text{-}C_{i+1}{}^b \equiv \delta^+\text{-}C_{x+1}{}^b)$ makes explicit the content-identity between α_{i+1} and $\delta_{x+1}{}^+$. We limit B_M+'s sensitivity to at least one instance of $\delta_{x+1}{}^+$ instead of the general set of δ_{x+1} to not run the risk of a sorites. Lastly, the content-identity relation involved here suffices to establish neither α-$C_{i+1}{}^b$ nor R* for the (α_i, α_{i+1}) case-pair, thereby safeguarding A+ as necessary for establishing R*.[17] These changes from B+ to B_M+ are reasonable given how much of a safer claim B_M+ is to B+, both in terms of sorites risk and the extent of what is demanded from the modal landscape in accommodating the principle. Note that B_M+ is meant to rationalise connecting α-$(BC)_i{}^a$ to α-$(BC)_{i+1}{}^b$ by offering up a proxy connection between α-$(BC)_i{}^a$ and δ^+-$(BC)_{x+1}{}^b$. However, we also want to convincingly account for α-$(\neg BC)_{i+2}{}^c$ to not instantiate a sorites. In other words, we *still* need to motivate the truth of α-$(BC)_{i+1}{}^b$ and not just δ^+-$(BC)_{x+1}{}^b$, given the desired establishment of R* through A+, while being sensitive to α-$(\neg BC)_{i+2}{}^c$, given the required avoidance of sorites.[18] As it stands, B_M+ and A+ jointly are not up to both tasks.

This connection, and thus R* itself, is borne out of a recognition that doxastic safety cannot be trivially guaranteed given our less-than-perfect discriminative

→ R* could simply be replaced with (A+ ∧ δ-$(BC)_x{}^b$) → R* and still be valid.

17 B_M+ follows closely Srinivasan's BEL* (2015, 302) and Vogel's (2010, 562) doxastic principle of 'indiscriminability-3', although Srinivasan and Vogel employ the condition of phenomenal duplication and relevant similarity, respectively, between cases instead of B_M+'s condition of content-identity.

18 This sensitivity also avoids Berker's (2008) criticism that modalized variants of otherwise soritical belief principles—e.g., variants of B+, such as B_M+—can lead "to unacceptable—or at least highly controversial—consequences" when iterated upon, such as there being "a possible case β such that *one feels as if one were in the center of the sun, and yet one believes that one feels cold*." (7n) This is because something like δ^+-$(BC)_n$ can be resisted if α-$(BC)_{n-1}$ never obtains in the first place.

capacities, and that, consequently, knowledge ends up being supported by rules entailing the propagation of belief over a range of cases satisfying doxastic safety for reasons that do not deal with such capacities. One way of thinking about this is that, *ceteris paribus*, satisfying safety principles is generally less demanding than satisfying margins principles.[19] This should be sensible since, one, safety principles usually incur a no-false-belief requirement while margins principles incur a no-falsity requirement for the range of cases over which these principles infer, and two, satisfaction of a no-falsity condition is more susceptible to the modal space being the way that it is than satisfaction only of a no-false-belief condition. This is because excluding the set of relevant falsehoods from a set of cases implies a larger extent of excluded instances than what obtains when just excluding the set of relevant false beliefs from the same set of cases—the set of falsehoods contains the set of false beliefs after all. This entails ostensibly *a priori* a greater difficulty in satisfying a no-falsity condition than a no-false-belief one, *ceteris paribus*. With this, it should come as no surprise that requiring only the satisfaction of safety for knowledge more easily accommodates solutions of trivial satisfaction via auspicious discriminative capacities than when knowledge also requires satisfying a margins principle, as while leveraging such capacities satisfies the former, these only go partway in satisfying the latter—hence the need for doxastic propagation principles like B_M+ and those that can connect δ^{+}-$(BC)_{x+1}{}^{b}$ to α-$(BC)_{i}{}^{a}$.

This consideration has three importantly related ramifications. First, it is not itself trivial that trivially safe beliefs are themselves trivially margins-satisfying, for one's discriminative capacities may lead to beliefs that are true but not factive over more cases than these true-belief ones. Second, beliefs can be safe even with inauspicious modal conditions, like when nearby cases are all (¬C)-verifying despite the actual case verifying C, if discriminative capacities are auspicious enough.[20] Third, these same modal conditions directly undermine beliefs being margins-satisfying since this necessitates that nearby relevant cases are *not* (¬C)-verifying. These ramifications should still be relevant to our discussion of A+ and R* despite their modal differences, with R* only being sensitive to actual cases while A+ is sensitive to general ones, because, as we justify below, the rules that derive R* also derive margins principles more similar in scope to R+ (see Section VIII)—i.e., those conceiving of knowledge as factive over a subset of N\PN-related general cases, not just neighbouring actual ones.

---

[19] Compare, for instance, Vogel's (2010) notions of strong and weak reliability (549, 566n12).

[20] This is reminiscent of Sosa's (2009) distinction between a "*brutely* reliable" epistemic competence and "the circumstances of the operation of that competence," where the former can be true whether the latter is "propitious" or not. (223n)

What anti-luminists like Williamson desire to account for by invoking phenomenal indiscrimin*ability* then is precisely a diminishment of the explanatory need for capacities that trivially ensure safety by way of making beliefs perfectly discriminative concerning the scope of cases comprehended by the safety principle at play. This indiscriminability principle, and the margins requirements it supports, makes knowledge more implausible than what mere regard for a no-false-belief condition of safety would have us suppose, because excluding the set of relevant false beliefs—what needs to occur for doxastic safety—cannot be ensured by appealing to maximally reliable belief formation borne out of perfect discriminative capacities that themselves are thrown into question. The same goes for the prior need of connecting $\delta^{+}\text{-}(BC)_{x+1}{}^{b}$ to $\alpha\text{-}(BC)_{i+1}{}^{b}$, for it appeals to a similar notion that motivates principles like B+: $\alpha\text{-}(BC)_{i}{}^{a}$ and $\alpha\text{-}(BC)_{i+1}{}^{b}$ are connected because our doxastic capacities are not discriminative enough to break that connection. The trick then is to construct a doxastic propagation principle *convincing enough* that appealing to the trivial safety of belief becomes no longer warranted as a realistic means of satisfying safety for knowledgeable belief.

B_{M}+ may constitute part of the picture of that principle, but its motivation has yet to be convincing. What will be shown, therefore, is that a natural way of deriving B_{M}+ through the idea of doxastic dispositions—i.e., dispositions to believe—is convincing in a manner that gives the anti-luminist unique conceptual resources for resisting the luminist in their charge that modes of knowledge undergirded by some form of doxastic privileged access can also defeat the relevance of epistemic margins principles. Starting with the next Section, we analyse both Srinivasan's (2015) understanding of doxastic dispositions as grounding principles of phenomenal indiscriminability and Vanrie's (2020) response to it to evaluate how B_{M}+ may fare as such a principle. This will also grant us the means to effectively fulfill the need of connecting $\delta^{+}\text{-}(BC)_{x+1}{}^{b}$ to $\alpha\text{-}(BC)_{i+1}{}^{b}$ while being sensitive to $\alpha\text{-}(\neg BC)_{i+2}{}^{c}$, which should help negate any risk of sorites along the $\alpha_{0\text{-}n}$ case-series.

IV

The idea that dispositions are what make B_{M}+ plausible ought to not surprise us. Given the desired avoidance of sorites, we want doxastic propagation principles that are also sensitive to when the propagation of beliefs is halted, but in a manner that neither begs the question against nor concedes too much to the luminist. Introducing the notion of a disposition to believe into the picture of the ALA is helpful insofar as it offers a good explanation for said sensitivity. One way of cashing this out is through an asymmetric relation between beliefs and dispositions in terms of what is reasonable to imply from one doxastic state to another. Broadly speaking, what we can say about beliefs given a present disposition to believe, and *vice versa*, may entail that within the set of N\PN-related indiscriminable cases in which you have the disposition to believe, some of those cases actualise that belief while the rest do not,[21] while a relation from beliefs to dispositions may support an entailment of differing strength and case jurisdiction.

Vanrie (2020) represents these thoughts by two principles. First is DISP-BEL, wherein, "[i]f in a case α, S has some disposition to believe that condition R obtains, then there is a sufficiently similar phenomenal duplicate β of α in which S believes that R obtains." (542) R here is a general proposition that can either be luminous or non-luminous, so we can rewrite the principle to be more in line with the language of B_M+ where, instead of phenomenal duplication, sufficient similarity, and possible cases (β), we have, respectively, content equivalency, N\PN-relation, and general cases (δ). If we write 'disposition to believe C' as '$B_{disp}C$', then we can rewrite DISP-BEL as,

$$\textbf{D-}B_M\textbf{+} \Leftrightarrow (\alpha\text{-}(B_{disp}C)_i{}^a \rightarrow [\delta^+\text{-}(BC)_x{}^a \wedge (\alpha\text{-}C_i{}^a \equiv \delta^+\text{-}C_x{}^a)]).$$

The second principle is DOXDIS, wherein, "[i]f in a case α, S believes that condition R obtains, then for any case β sufficiently similar to α, S has some disposition in β to believe that R obtains." (541)[22] Relating once again to B_M+, we can rewrite DOXDIS as,

$$\textbf{DOX-}B_M\textbf{+} \Leftrightarrow (\alpha\text{-}(BC)_i{}^a \rightarrow \delta\text{-}(B_{disp}C)_x{}^b),$$

whereby cases of belief infer cases of doxastic disposition in any N\PN-related indiscriminable case.[23]

D-B_M+ and DOX-B_M+ not only specify a *mutual* relation between beliefs and their underlying dispositions, which is important in starting to flesh out the asymmetric belief/disposition relation,[24] but they also jointly derive B_M+.[25] This derivation is analogous to how Vanrie (2020, 542-543) derives BEL* from DOXDIS

[21] Srinivasan (2015, 303n19) makes a similar observation.

[22] DOXDIS is adapted from Srinivasan's (2015, 303) own formulation of the same principle.

[23] Note that we do not suppose that $\{\alpha_i\} \in \{\delta_x\}$ in order to avoid the inference, $\alpha\text{-}(BC)_i{}^a \rightarrow \alpha\text{-}(B_{disp}C)_i{}^a$, because this naively denies the possibility of one believing in something without a corresponding disposition. Additionally, D-B_M+ and DOX-B_M+ can be generalised further to include δ in the antecedent and not just α, but the rest of the discussion going forward does not turn on this inclusion.

[24] The plausibility of D-B_M+ and DOX-B_M+ can also be independently motivated through what they imply about the modal space. See Note 40 and the context provided in Section VII.

[25] *Proof*: Given DOX-B_M+, we can deduce $\alpha\text{-}(BC)_i{}^a \rightarrow \alpha\text{-}(B_{disp}C)_{i+1}{}^b$ due to $\{\alpha_{i+1}\} \in \{\delta_x\}$. Given D-$B_M$+, and plugging in $\alpha\text{-}(B_{disp}C)_{i+1}{}^b$ into its antecedent, we get $\alpha\text{-}(BC)_i{}^a \rightarrow [\delta^+\text{-}(BC)_{x+1}{}^b \wedge (\alpha\text{-}C_{i+1}{}^b \equiv \delta^+\text{-}C_{x+1}{}^b)]$.

and DISP-BEL. However, he also contends that establishing BEL* in this way can result in a sorites depending on how one understands the relation between BEL*-required cases that are phenomenal duplicates and those that are sufficiently similar to each other. (543) For B_{M}+, this would be the relation between cases that have identical content and those that are mutually N\PN-related and indiscriminable e.g.—between the case-pair (α_i, δ_x^+) and the pair (α_i, α_{i+1}). In fact, we can indeed establish B_{M}+ as soritical in this manner: consider a rule, DUPL-N/I-1, whereby any case, X, N\PN-related to and indiscriminable from another case, Y, is also an N\PN-related and indiscriminable case to any case, Z, whose content is identical to Y;[26] this permits α_{i+1} being an N\PN-related and indiscriminable case to δ_x^+, meaning that, by D-B_{M}+ and then DOX-B_{M}+, $\alpha\text{-}(B_{disp}C)_{i+1}{}^{b}$ is true given $\alpha\text{-}(B_{disp}C)_i{}^{a}$. This concludes a sorites since iterating on this process grants that same disposition for S in case α_n, and it is quite uncontentious that S must have possessed a disposition to believe that C at some point in the ALA's case series anyways. Moreover, the disposition actualising as $\alpha\text{-}(BC)_n$ would contradict the "incontestable [notion] that BC does not obtain in case α_n" (Berker 2008, 8).

Obviously, the above sorites can be halted by having N\PN-relation and indiscriminability *fully* transitive over content identity instead of, as DUPL-N/I-1 would have it, merely *partly* transitive over it. Consider the rule, DUPL-N/I-2, whereby any case, X, N\PN-related and indiscriminable concerning another case, Y, is also an N\PN-related and indiscriminable case to any case, Z, whose content is identical to Y, *if Z's content is also identical to that of X*. This prevents (α_{i+1}, δ_x^+) from necessarily being a pair of N\PN-related and indiscriminable cases, thereby preventing the trivial inference to $\alpha\text{-}(B_{disp}C)_{i+1}{}^{b}$.[27] DUPL-N/I-2 is much more restrictive than DUPL-N/I-1 in specifying what suffices for determining (α_{i+1}, δ_x^+) as a pair of N\PN-related, indiscriminable cases. Accordingly, if one had to appeal to either of the two rules, then while one would not need X and Y to be content-identical for Z and X to be N\PN-related, only through DUPL-N/I-2 would there be no guarantee of such N\PN-relation outside of X and Y being content-identical—if Y is content-identical to Z which is content-identical to X, then Y and X are content-identical pairs—which, in the case of (α_i, α_{i+1}) *necessarily* fixes their content relation *m*-value to zero if one desires for either (α_{i+1}, δ_x^+) or (α_i, δ_{x+1}^+) to be *guaranteed* as N\PN-related case-pairs.

This is all relevant to the anti-luminist, for one way of connecting $\delta^+\text{-}(BC)_{x+1}{}^{b}$ to $\alpha\text{-}(BC)_{i+1}{}^{b}$ in order to derive R* is by having something like $\delta^+\text{-}(B_{disp}C)_{x+1}{}^{b}$ feature

[26] See Vanrie (2020, 543) for the analogous rule, DUPL-SIM-1.

[27] An analogous attempt at halting the sorites is given in Vanrie (2020, 545) via application of the rule, DUPL-SIM-2.

in sufficient conditions for $\alpha\text{-}(BC)_{i+1}{}^{b}$, with the former being derivable from $\alpha\text{-}(BC)_{i}{}^{a}$ via DOX-B_M+, but only if $\{\delta_{x+1}{}^{+}\} \in \{\delta_x\}$. In effect, the desired connection first requires $\delta_{x+1}{}^{+}$ and α_i to be N\PN-related, but guaranteeing this through DUPL-N/I-1 leads to a sorites, as argued above, and otherwise through DUPL-N/I-2 one is left with overly strict requirements. A plausible rule would need to avoid both issues, preferably by appealing to case features that are neither trivially guaranteed nor otherwise too restrictive on the relevant cases.

A plausible candidate approach, called DUPL-N/I-3, has it that the N\PN-relation is *fully* transitive over sufficient similarity—e.g., $\delta_{x+1}{}^{+}$ and α_i are N\PN-related if(α_i, α_{i+1}) and (α_{i+1}, $\delta_{x+1}{}^{+}$) are sufficiently similar pairs of cases. In general, DUPL-N/I-3 has it that any case, X, sufficiently similar to another case, Y, is N\PN-related to any case, Z, sufficiently similar to Y. DUPL-N/I-3 and DOX-B_M+ jointly derive $\delta^{+}\text{-}(B_{disp}C)_{x+1}{}^{b}$ in all instances where (α_{i+1}, $\delta_{x+1}{}^{+}$) is a sufficiently similar case-pair since the ALA is set up to have (α_i, α_{i+1}) be sufficiently similar already.[28] This avoids the sorites occurrent from DUPL-N/I-1 because (α_n, $\delta_{n-1}{}^{+}$), which would need to be an N\PN-related and indiscriminable case-pair for the sorites to go through, is not necessarily so since neither (α_{n-1}, $\delta_{n-1}{}^{+}$) nor (α_n, $\delta_n{}^{+}$) are necessarily sufficiently similar pairs: content-identical cases are not also necessarily sufficiently similar if there is even the slightest possibility that differences between them that have nothing to do with their content—i.e., differences in how cases are specified that do not involve the propositions in question—are extensive enough to warrant them not being sufficiently similar. Obviously, this needs to be a very nonproblematic possibility lest we risk conflating cases with any one of their contents, but this possibility ought to go through without much resistance since cases are distinct from propositions and are not specified wholly by propositional content.

All that is left is to connect $\delta^{+}\text{-}(B_{disp}C)_{x+1}{}^{b}$ to $\alpha\text{-}(BC)_{i+1}{}^{b}$. Let us have it that actualised doxastic disposition—i.e., disposition that actualises as a belief—infers belief in all cases content-identical and sufficiently similar to the case of actualised disposition: if we write, 'B is sufficiently similar to A', as, 'A $\Rightarrow_{SS}$ B', then, $\forall[(a, \delta, i, x^*) : (\delta_i \Rightarrow_{SS} \delta_{x^*})]$,

$$\text{DD-}B_M\text{+} \Leftrightarrow ([\delta\text{-}(BC)_i{}^{a} \wedge \delta\text{-}(B_{disp}C)_i{}^{a}] \rightarrow \delta\text{-}(BC)_{x^*}{}^{a}).$$

'x*' is meant to label any general case (δ), actual (α) or possible (β), that is sufficiently similar to δ_i. There is no absolute restriction on the numerical value for 'x*', just like

[28] This implies sufficient similarity also for the (α_{i+1}, α_{i+2}) case-pair, which is meant to contrast, in Section V, with how the ($\delta_{x+1}{}^{+}$, $\delta_{x+2}{}^{+}$) case-pair can concomitantly get away with being insufficiently similar given how the $\alpha_{0\text{-}n}$ case-series is not subject to the same modal restrictions as the one involving $\delta_x{}^{+}$.

for ‘x’, but generally, $\forall$(i, x*, y, z), ($0 \leq i \leq n$), ($i\text{-}y \leq x^* \leq i+z$), ($1 \leq y \leq i+n+1$), ($1 \leq z \leq n\text{-}i+1$). Now, DD-$B_M$+ does not suffice for α-$(BC)_{i+1}{}^{b}$ since, given B_M+, it is not automatically the case that ($\delta_{x+1}{}^{+} = \delta_{x^*+1}{}^{+}$). In other words, the BC-verifying ($\delta_{x+1}{}^{+}$)-case that is the consequent of B_M+ needs to be sufficiently similar to α_{i+1} before, through DUPL-N/I-3, DOX-B_M+, DD-B_M+, and then A+, R* is derived. As such, DD-B_M+ helps derive R* in a qualified sense:

$\forall$[(α, δ, i, x, x*) : ($\delta_i \Rightarrow_{SS} \delta_{x^*}$), ($\alpha_i \Rightarrow_{SS} \alpha_{i+1}$), ($\alpha_i \Leftrightarrow_{N\backslash PN} \delta_x$)], $\exists$($\delta_{x^{*+}}$) : ($\{\delta_{x^{*+}}\} \subset \{\delta_x\}$), ($\delta_{x+1}{}^{+} = \delta_{x^*+1}{}^{+}$), [**($B_M$+, DUPL-N/I-3, DOX-$B_M$+, DD-$B_M$+, A+) $\rightarrow$ R***].

Relatedly, DD-B_M+ would still not suffice for R* even if, instead of inferring BC in cases sufficiently similar to a (BC $\wedge$ B_{disp}C)-verifying ($\delta_{x+1}{}^{+}$)-case, it inferred BC in cases N\PN-related to it, because the BC-verifying ($\delta_{x+1}{}^{+}$)-case inferred by B_M+ is not guaranteed, by any of the rules or set-up for the ALA, and even given its N\PN-relation to α_{i+1}, to even be B_{disp}C-verifying since it is necessarily N\PN-related to neither α_i nor some BC-verifying ($\delta_x{}^{+}$)-case. Additionally, rewriting DD-B_M+ in this fashion attains problematic ramifications. One, most pressingly, to prevent a sorites (absurdly concluding BC for α_n), such a ($\delta_{x+1}{}^{+}$)-case cannot be B_{disp}C-verifying, meaning that, because of DOX-B_M+ and DUPL-N/I-3, it can never be sufficiently similar to α_{i+1}. This is too strong a restriction on what case-pairs can be sufficiently similar since, as discussed below, preventing a sorites can occur by just applying this restriction to (α_{i+2}, $\delta_{x+2}{}^{+}$) case-pairs, not also to (α_{i+1}, $\delta_{x+1}{}^{+}$) ones. Two, DD-B_M+ rewritten in this way is itself unnecessarily permissive, because DD-B_M+ can already halt a sorites without needing to infer belief in the relatively larger set of N\PN-related cases.

Indeed, DD-B_M+ as currently conceived can even accommodate (α-$(KC)_i{}^{a}$ $\wedge$ α-$(\neg BC)_{i+1}{}^{b}$) as an (R*)-consistent result, as it is not impossible under the right modal circumstances for, due to ($\delta_i \Rightarrow_{SS} \delta_{x^*}$) and either ($\delta_{x+1}{}^{+} \neq \delta_{i+1}{}^{+}$) and/or ($\alpha_{i+1} \neq \delta_{x^*+1}{}^{+}$), no ($B_{disp}$C $\wedge$ BC)-verifying ($\delta_{x+1}{}^{+}$)-case to be sufficiently similar to α_{i+1}. How this (R*)-consistency is explained can also explain how DD-B_M+ precludes a sorites: cases of actualised disposition may be far-enough possibilities in relation to the actual case, perhaps influenced by a very weak disposition in the actual case to believe that C, that no risk obtains of the immediately adjacent case also being BC-verifying. This ought to be favourable to anti-luminists on account of the scope of what a margins principle is meant to require. Such principles, at least when used to analyse the ALA, demand only factivity for knowledge over a range of cases, not that knowledge infer *belief* over this same case-range as well.[29]

[29] Srinivasan (2015) also indicates how dispositions can be usefully employed to mark the threshold at which beliefs start to fall off. However, as shown by Vanrie (2020), one can deny the sense

One other positive motivation for DD-B_{M+} is that it offers natural interpretations of beliefs and dispositions to believe. If, by DOX-B_{M+}, we accept that cases of belief infer cases of doxastic disposition in any N\PN-related and indiscriminable case, then those dispositions must, to be regarded as dispositions at all, actualise somehow. D-B_{M+} is as strong as needed to help deduce B_{M+}, but α-$(BC)_{i+1}{}^{b}$ requires something stronger to infer it from δ^{+}-$(BC)_{x+1}{}^{b}$. DD-B_{M+} contributes its strength for the inference, but having this strength be insufficient for the task not only helps resist sorites but does so in a manner that lends credence to the notion that what explains an absence of belief often has to do with factors that have nothing to do with us or even with rules of doxastic propagation—e.g., B_{M+} and DOX-B_{M+}—that characterise our relatively indiscriminative doxastic capacities. Additionally, DD-B_{M+} respects the intuitive claim that beliefs should not start to fall off *only* when cases become phenomenally discriminable, for then we would be forced to espouse a guarantee of belief in all sufficiently similar and *phenomenally indiscriminable* cases to the case of actualised disposition, which is a much stronger condition than DD-B_{M+} alone and would also be conceding too much to the luminist.

Lastly, we can compare DD-B_{M+} with a relatively close variant, DISP-BEL-EXTREME, which Vanrie (2020) considers to be an *ad hoc* rule due to its licensing the inference from our dispositions to not just nearby belief but "*extremely* nearby belief." (546) DISP-BEL-EXTREME, when adapted to the language of N\PN-relation and content identity, is essentially a reworking of D-B_{M+}, wherein instead of ($\alpha_i \Leftrightarrow_{N\backslash PN} \delta_x$) we have, ($\alpha_i \Rightarrow_{ES} \delta_x$), where 'ES' stands for 'extreme similarity'. Extreme similarity is a simple relation over which sufficient similarity is fully transitive.[30] Nevertheless, DD-B_{M+} should not be that much more implausible than DOX-B_{M+}, for while the latter comprehends an inference from belief, the former comprehends an inference from dispositions. Moreover, even if we assume that deriving dispositions from belief is easier than deriving belief from dispositions, DD-B_{M+} requires only a more limited modal space for specifying what is inferable from beliefs and dispositions than what DOX-B_{M+} employs for inference from beliefs—i.e., what can be said about actualised dispositions among sufficiently similar and content identical cases versus what can be said about beliefs among the larger space of N\PN-related and indiscriminable cases—especially since sufficient similarity and N\PN-relation are connected via NeighSim. This implies that labelling DD-B_{M+} as *ad hoc*

ascribed to dispositions that make them function in this way without being committed to problematic ways of relating beliefs and dispositions. I aim to show below another way of ascribing a sense to dispositions, especially in how it relates to safety, that does force such a commitment if denied.

[30] See Vanrie (2020, 546) for the rule, EXTREME-TRANS.

contributes to delegitimising DOX-B_{M}+ as well. Importantly, as is explored below, the anti-luminist can employ both DD-B_{M}+ and DOX-B_{M}+, or even principles like them, for significant dialectical advantage over the luminist's anti-margins stance.

V

Recall that we need to connect δ^{+}-$(BC)_{x+1}{}^{b}$ to α-$(BC)_{i+1}{}^{b}$ while being sensitive to α-$(\neg BC)_{i+2}{}^{c}$ in order to establish R* without succumbing to a sorites involving belief and/or disposition propagation throughout a case-series. Doing so can be accomplished by accounting for α-$(\neg C \wedge \neg BC)_{i+2}{}^{c}$, regardless of α_{i+1} being BC- or ($\neg$BC)-verifying. For the sake of simplicity, this Section onwards, unless otherwise indicated, will focus on α-$(BC)_{i+1}{}^{b}$ being true, for the anti-luminist's resistance of luminism does not turn on a necessitation of α-$(BC)_{i+1}{}^{b}$.[31] In effect, the anti-luminist gains an advantage over the luminist by their account being more motivated than the luminist's. How this works boils down to, one, preventing α-$(KC)_{i+1}{}^{b}$ by having at least one (BC $\wedge$ $\neg$C)-verifying (δ_{x+2})-case be N\PN-related to and indiscriminable from α_{i+1}—i.e., one verifying δ-$(BC \wedge \neg C)_{x+2}$—and two, having no such (BC $\wedge$ $\neg$C)-verifying (δ_{x+2})-case be both sufficiently similar and content-identical to α_{i+2}. The first condition precludes the luminist's desired trivial verification of A+ for (α_{i+1}, δ_{x+2}) case-pairs, while the second precludes a trivial falsification of α-$(\neg BC)_{i+2}{}^{c}$ sourced from the application of DUPL-N/I-3, DOX-B_{M}+, and DD-B_{M}+ given α-$(BC)_{i+1}{}^{b}$. Meeting both conditions through reasonable means implies the anti-luminist's ability to account for α-$(\neg BC)_{i+2}{}^{c}$ without needing to invoke the type of explanatory basis the luminist would want to halt sorites with, that being an auspicious enough discriminative capacity.

To see how the anti-luminist comes out on top, first note that all that α-$(\neg BC)_{i+2}{}^{c}$ requires is that, given DD-B_{M}+, no (BC $\wedge$ B_{disp}C)-verifying case is both sufficiently similar and content-identical to α_{i+2}. With the additional requirement of not trivially verifying A+ for (α_{i+1}, δ_{x+2}) case-pairs, we have several overlapping results permitted by both requirements. However, there are some results that satisfy or are allowed by one while falsifying the other. For instance, if a (BC $\wedge$ $\neg$C)-verifying (δ_{x+2})-case is sufficiently similar and content-identical to α_{i+2}, then while trivial verification of A+ is stopped, α-$(\neg BC)_{i+2}{}^{c}$ becomes falsified due to DUPL-N/I-3, DOX-B_{M}+ given α-$(BC)_{i+1}{}^{b}$, and DD-B_{M}+. The converse outcome—keeping α-$(\neg BC)_{i+2}{}^{c}$ while ensuring trivial verification of A+ for (α_{i+1}, δ_{x+2}) case-pairs—can be met by having (BC $\wedge$ $\neg$C)-verifying (δ_{x+2})-cases be neither, one, both sufficiently similar and content-identical to α_{i+2}, nor, two, both N\PN-related to and

[31] Note 36, along with the fuller context provided by Section VII, substantiates this.

indiscriminable from α_{i+1}. This first outcome is undesirable by both luminists and anti-luminists given the set-up of the ALA while the second one is desirable only to the luminist as it permits a result like (α-$(KC)_{i+1}{}^{b}$ ∧ α-$(\neg C)_{i+2}{}^{c}$), which contradicts R*. Consequently, the anti-luminist must avoid both outcomes. The only way to achieve this then would be if the only (BC ∧ ¬C)-verifying (δ_{x+2})-cases are those that can be both N\PN-related to and indiscriminable from α_{i+1} but not both sufficiently similar and content-identical to α_{i+2}. The anti-luminist and the luminist therefore differ in their evaluation of (δ_{x+2})-cases in terms of the result that (BC ∧ ¬C)-verifying cases can be both N\PN-related to and indiscriminable from α_{i+1}. If the anti-luminist's motivation for this result is better than the luminist's motivation for its negation, then the anti-luminist will come out on top.

A primary attempt at doing so leverages the fact that there is nothing in the rules establishing R* that precludes (δ_{x+2})-cases N\PN-related to and indiscriminable from α_{i+1} also being (BC ∧ ¬C)-verifying. Yes, B_M+ requires, given α-$(BC)_{i+1}{}^{b}$, at least one BC-verifying (δ_{x+2^+})-case that is N\PN-related to α_{i+2}, *not* α_{i+1}, but nothing about B_M+ prevents at least one (BC ∧ ¬C)-verifying case from the broader set of (δ_{x+2})-cases from being N\PN-related to both. Moreover, there is also nothing preventing these cases from being content-identical to, or just indiscriminable from, α_{i+2}. All that is required, after all, for R* to not risk sorites is for these cases to not also be sufficiently similar to α_{i+2}. Nevertheless, seeking motivation by mentioning mere possibility is not an optimal path of substantiation since the possibility can end up being too weak for the anti-luminist's benefit.

I argue that a better way for the anti-luminist would be to leverage an already excavatable distinction between beliefs and dispositions to believe by which δ-$(BC \wedge \neg C)_{x+2}$ can obtain without negating α-$(\neg BC)_{i+2}{}^{c}$ nor trivialising A+. This is because such a distinction naturally comes out from the rules that establish R* in a manner that negates not only sorites but A+'s trivialisation as well, and denying said distinction to safeguard A+'s trivialisation leads to highly unintuitive luminist commitments about how sorites is halted. The general argument is outlined in four parts:

> P1) $B_{disp}C$-verifying (δ_{x+2})-cases can be both sufficiently similar and content-identical to α_{i+2} while BC-verifying ones, if content-identical, cannot also be sufficiently similar to it.
>
> P2) There can be differences in the truth-values for BC and $B_{disp}C$ across case-pairs that have identically indiscriminable content differences—e.g., a (α_{i+1}, δ_{x+2^+}) case-pair can represent a loss of $B_{disp}C$ but not BC (i.e., mere disposition-loss given by δ^+-$(\neg B_{disp}C \wedge BC)_{x+2}{}^{c}$) while the ($\alpha_{i+1}$, α_{i+2}) pair represents a loss of BC but not $B_{disp}C$ (i.e., mere belief-loss given by α-$(B_{disp}C \wedge \neg BC)_{i+2}{}^{c}$) even though ($\alpha_{i+1}$, δ_{x+2^+}) and (α_{i+1}, α_{i+2}) are identically indiscriminable (i.e., share the same *m* value).

P3) P2 can be explained by introducing a novel discriminative dimension, termed 'dispositional discriminability', that can be substantiated from the rules establishing R* while also substantiating both P1 and the non-trivialisation of A+.

P4) Dispositional discriminability is hard to deny without committing to highly unintuitive assumptions.

P1 allows for α-$(\neg BC)_{i+2}{}^{c}$ without contradicting the rules establishing R*. P2 follows from P1 and is motivated by its mere possibility opening the path towards an explanation, in P3, that not only favours anti-luminism but does so in a way, via P4, that strengthens the case for A+'s non-trivialisation. I address P1 and P2 for the rest of this Section, P3 in Sections VI and VII, and P4 from Section VII onwards.

Indeed, if we take α-$(BC \wedge B_{disp}C)_{i+1}{}^{b}$ to represent a case N\PN-related to and indiscriminable from α_{i+2}, then by DOX-B_M+, α_{i+2} must be both $(\neg BC)$- *and* $B_{disp}C$-verifying. This denotes an instance of a general characterisation of the belief/disposition distinction being argued for here: beliefs and dispositions do not necessarily rise and fall together along a case-series. This rationalises P1 – that $B_{disp}C$-verifying (δ_{x+2})-cases are closer to α_{i+2} than BC-verifying ones are—not only because of concordance between P1 and α_{i+2} itself being $B_{disp}C$-verifying, but also because it is a permitted outcome by the R*-establishing rules.

Nonetheless, what is also permitted by said rules is the possibility of some $(\neg B_{disp}C)$-verifying (δ_{x+2})-cases that are modally close to α_{i+2} but not, if also BC-verifying, both content-identical and sufficiently similar to it. This is recommended for strengthening the cause for α-$(\neg BC)_{i+2}{}^{c}$ without contradicting it through α-$(BC \wedge B_{disp}C)_{i+1}{}^{b}$ and DOX-B_M+, because otherwise it becomes open for δ-$(B_{disp}C)_{x+2}$ to be a frequent enough result within modal space to where claiming that α-$(\neg BC)_{i+2}{}^{c}$ obtains *rather than* α-$(BC)_{i+2}{}^{c}$ becomes less reasonable.

As such, along with the result that any $B_{disp}C$-verifying (δ_{x+2})-case sufficiently similar and content-identical to α_{i+2} cannot be BC-verifying, we have it that within the set of (δ_{x+2})-cases content-identical to α_{i+2} we can find $B_{disp}C$-, BC-, $(\neg B_{disp}C)$-, and $(\neg BC)$-verifying ones, albeit with $B_{disp}C$-verifying ones being closer to α_{i+2} than BC-verifying ones.[32] This speaks to two implications: one, that differences in beliefs and dispositions between cases can result in identical content between them, plausibly because case content matters for the truth of C, and neither beliefs nor dispositions strike as what would directly matter to propositional truth given subject-matter differences between them; and two, this variety of results can be explained by the fact that having a doxastic disposition without its corresponding actualisation as a

[32] Even if we assume that all $(\delta_{x+1}{}^{+}, \delta_{x+2}{}^{+})$ case-pairs are N\PN-related, the reason why $(\delta_{x+2}{}^{+})$-cases can be $(\neg B_{disp}C)$-verifying is because not every $(\delta_{x+1}{}^{+})$-case has to be BC-verifying if B_M+ only requires one of the latter.

belief is perfectly understandable according to the disposition *as* a disposition to believe, while having a belief without a corresponding and underlying disposition strikes as a rarer situation altogether. The first implication is a straightforward inference from the above discussion. The second implication is not as easily inferable as it requires a careful analysis of $B_{disp}C$-, BC-, and ($BC \land B_{disp}C$)-verifying cases and their modal/discriminative relations to α_{i+1} and α_{i+2}. Both implications are more fully discussed in the next two Sections.

Regardless, for our purposes now, it suffices to comment that the picture that surfaces here is one of beliefs and dispositions falling off at different times, expressible as differences between α_{i+2} and δ_{x+2}^{+} that matter to the combination of DOX-B_{M}+, D-B_{M}+, and DD-B_{M}+ not leading to a sorites in either beliefs or dispositions propagating along a case-series. These differences obtain because the series involving δ_{x}^{+} is one of a subset of *general* cases – i.e., it represents a subset of cases that exhibit differences in mutual relations, such as (δ_{x+1}^{+}, δ_{x+2}^{+}) being either a sufficiently similar case-pair (representing mere belief-loss) or even a non-N\PN-related one (representing mere disposition-loss), that are not possible for the case-series, $\alpha_{0\text{-}n}$, given its set-up consisting of mutually sufficiently similar pairs of neighbouring cases. These differences are not arbitrarily chosen since they are consistent with the rules already establishing R*.

Importantly, this result of greater relational freedom of $B_{disp}C$-verifying (δ_{x+2}^{+})-cases compared to BC-verifying ones, which is required to safeguard α-$(\neg BC)_{i+2}{}^{c}$, favours the anti-luminist position over the luminist's own. The luminist's position would be that δ^{+}-$(BC)_{x+2}{}^{c}$ not obtaining for sufficiently similar (α_{i+2}, δ_{x+2}^{+}) case-pairs is because of one's auspicious discriminative capacities, while the anti-luminist's position would instead have to appeal to the modal landscape for this result. However, these explanations are not equivalently demanding: all the anti-luminist needs to guarantee is for the result of δ^{+}-$(BC)_{x+2}{}^{c}$ not obtaining for sufficiently similar (α_{i+2}, δ_{x+2}^{+}) case-pairs to happen *at least once* in the case-series for their establishment of R* to be sensitive to α-$(\neg BC)_{i+2}{}^{c}$, while the luminist must guarantee a discriminative capacity that can apply at *any* point in the case-series by which BC would otherwise be unsafe. The luminist's guarantee obviously requires more than the anti-luminist's and thus convincingly supports the anti-luminist's attempt at permitting α-$(\neg BC)_{i+2}{}^{c}$. Note the nuance of the anti-luminist's position here. They expressly deny the luminist's move of appealing to an auspicious discriminative capacity because it is not this that, for the anti-luminist, avoids sorites. If one's discriminative capacity cannot do the job alone, then obviously the modal landscape can step in. The fact that the above combination of doxastic states that must happen for α-$(\neg BC)_{i+2}{}^{c}$ to not be contravened needs to obtain only at least once actually works in the anti-luminist's favour, as it all the more serves as reminder that, *precisely because* of the phenomenon of indiscriminability, results like $\alpha\text{-}(\neg BC \wedge \neg C)_{i+2}{}^{c}$ are not as possible as luminists might think.

Nonetheless, what this belief/disposition distinction means for A+'s non-trivialisation still needs fleshing out in a manner that does not beg the question against the luminist. Part of the way has already been paved above by conceiving the luminist position as more demanding than the anti-luminist's own. Still, this way has involved a distinction that demands explanation since the distinction can apply over a set of phenomenally *in*discriminable cases—i.e., the distinction must be explained by invoking something *other* than phenomenal indiscriminability that remains relevant for the contents with which one's beliefs and dispositions happen to be associated. Regardless of the terms in which this explanation ends up cashed out, what is apparent from the way the argument thus far has been set up is that luminists must deny this belief/disposition distinction lest they remain consistent with the way R* has been established to be sensitive to $\alpha\text{-}(\neg BC)_{i+2}{}^{c}$. If the anti-luminist can successfully defend this distinction by said explanation, then what the explanation implies about beliefs like $\alpha\text{-}(BC)_{i+1}{}^{b}$ being unsafe would thereby undermine the luminist's desire to have $\alpha\text{-}(BC)_{i+1}{}^{b}$ be safe and R* be false.

Williamson (2000) argues that the ALA not risking a sorites is based on our doxastic capacities being the way that they are. (103-104) I expand upon this by arguing that the above modal account of how the disposition-belief connection comes apart *is how one avoids a sorites* because the account expresses a plausible story of how knowledgeable belief and propositional falsity are distanced from each other in (R*)-consistent ways. As the story goes, why one cannot know that something is true when in an adjacent case it is false has to do with our dispositions being able to not only fall off but fall off at different points from when our beliefs fall off and when falsity arises. This is because to say that one's beliefs and doxastic dispositions rise and fall together all too easily permits, as fleshed out more in Section VI, the luminist's claim that (¬C)-verifying cases are simply those in which one could never falsely believe that C *due to not even having a disposition to believe that C*. This begs the question against Williamson by trivially verifying safety for our beliefs.

In summary, a more complete analysis of this belief/disposition distinction in P3 must capture three things. One, to not trivialize A+, it must permit both ¬C and BC for some δ_{x+2} that is either not sufficiently similar or not content identical to α_{i+2} yet still N\PN-related to and indiscriminable from α_{i+1}. Two, for another such δ_{x+2}, either the same case or another one, said distinction must also permit BC and $B_{disp}C$. We have already seen one possible example thereof, $\delta^{+}\text{-}(B_{disp}C \wedge BC)_{x+2}{}^{c}$ for insufficiently similar (α_{i+2}, $\delta_{x+2}{}^{+}$) case-pairs, but this ought to generalise to at least some (δ_{x+2})-cases, not just because otherwise one risks conceding to the luminist that ($\neg B_{disp}C \land \neg BC$) obtains for every other such (δ_{x+2})-case that is not content-identical to α_{i+2}, but also because otherwise the case-pairs to which DD-B_M+ is irrelevant arbitrarily become those to which DOX-B_M+ and D-B_M+ are also irrelevant. Lastly, three, among (δ_{x+2}^{+})-cases, $B_{disp}C$-verifying ones are the closest to α_{i+2} both modally and discriminatively.

VI

To analyse this distinction, we can take our cues from DD-B_M+ and DOX-B_M+ given that both involve *general* cases of belief and disposition. To remain faithful to the analysis thus far, the distinction ought to consist of both discriminative and modal elements. First, note that necessary conditions for DD-B_M+ and DOX-B_M+ are related: there is a phenomenal requirement of content-identity and indiscriminability for the former and latter, respectively, and a modal requirement of sufficient similarity and N\PN-relation for the former and latter, respectively. However, content-identity entails indiscriminability while sufficient similarity, by meeting some threshold of similarity (*s*) between cases, entails them being mutually N\PN-related. Therefore, there are stricter phenomenal and modal conditions for belief-propagation along a case-series than for disposition-propagation.

Nonetheless, this strictness is sensical given how a disposition without belief is intuitively a more common affair than a belief without an underlying disposition, because it is seemingly more understandable how a belief could function as evidence for an underlying disposition than a disposition evincing a belief, thereby making it harder to break the evidential implication of the former than that of the latter. Indeed, we can employ this observation to explain several pre-existing relations between DD-B_M+, D-B_M+ and DOX-B_M+. First, it explains why dispositions alone can only infer the *existence* of a belief in a relatively smaller sub-set of cases content-identical to the disposition-case (D-B_M+), while beliefs alone can infer a *general* set of dispositions in the relatively larger sub-set of cases indiscriminable to the belief-case(DOX-B_M+). Second, as hinted at in the end of Section IV, it explains how DD-B_M+ requires only a more limited modal and phenomenal space (sufficiently similar and content-identical cases) for specifying how beliefs are inferable from other cases of beliefs and dispositions than what DOX-B_M+ employs for inferring dispositions from beliefs (N\PN-related and indiscriminable cases), especially since sufficient similarity and N\PN-relation are connected via NeighSim.

Lastly, by comparing D-B_M+ with DD-B_M+, we can see an additional pattern: one can infer only the *existence* of a belief-case from that of a disposition-case (D-B_M+) while one can infer a *general set* of belief-cases from the existence of even a rarer kind of case that verifies both a disposition and a belief (DD-B_M+). A natural way of explaining this could be that, although dispositions are more easily inferable from beliefs than *vice versa*, beliefs are best inferred from other beliefs, not dispositions. We can also compare the broader phenomenal requirements of B_M+to those of D-B_M+ to reach the same conclusion comparing these two.

Importantly, and this is the crux of this Section, we can extrapolate from these patterns to create an analogue of DD-B_M+, called DDOX-B_M+, that infers general disposition-cases instead of belief-cases from cases that verify both dispositions and beliefs. DDOX-B_M+ thus functions as the final piece in our analysis of the belief/disposition distinction. Now, DDOX-B_M+ would need to specify combined modal/discriminative requirements of a greater extension than those of DD-B_M+ and DOX-B_M+ since, one, it is easier to infer dispositions than beliefs, *ceteris paribus*, and two, intuitively, cases of both disposition and belief function as stronger evidence for disposition-cases elsewhere compared to this same evidential implication from mere belief. Specifically, to not make DDOX-B_M+ too strong of a claim, the safest way to realise these constraints would be to have a general case verifying both BC and B_{disp}C also infer B_{disp}C in cases N\PN-related to *but not* indiscriminable from said general cases. This extension of the discriminative, but not modal dimension is plausible on account of N\PN-relation already being quite a broad relation to begin with. Not only can the N\PN-relation apply to far-off cases insufficiently similar to α_i, but more of the set of cases indiscriminable from α_i are within the set of cases N\PN-related to α_i than *vice versa*, for it is ostensibly more plausible for an N\PN-related case to also be discriminable from α_i than for an indiscriminable case to not be N\PN-related to it.[33]

Of course, DDOX-B_M+ does not entail that B_{disp}C can be found in all cases *discriminable* from and N\PN-related to a (BC ∧ B_{disp}C)-verifying one, since that would imply that one would have a disposition to believe that C even in a (¬C)-verifying case wherein the difference between the content in this case, mattering for whether C or ¬C obtains, and the content in a possibly adjacent C-verifying case is highly discriminable, which is absurd. Indeed, the use of indiscriminable so far has implied *phenomenal* indiscriminability, so we could cash out some conception of non-phenomenal discriminability such that a B_{disp}C-verifying case can be both phenomenally discriminable yet non-phenomenally indiscriminable from α_i. This could help halt the propagation of doxastic dispositions to cases highly

[33] Given NeighSim, one can hold to this result even while accepting the more luminist position that, as Berker (2008) does, "phenomenologically indistinguishable bases of belief need not count as sufficiently similar" (9n15), as it is less likely that such belief bases are also not N\PN-related due to N\PN-relation being broader than sufficient similarity.

phenomenally discriminable from α_i if they are also non-phenomenally discriminable.

The rationale for this is simple: beliefs ought to start falling off in cases modally closer to the initial (BC ∧ B_{disp}C)-verifying case than those in which dispositions start to fall off, because losing a disposition to believe is more of a reason to not have a belief than not believing is a reason to not have a disposition at all. This also ensures that, while BC-verifying cases guarantee B_{disp}C in a subset of cases phenomenally indiscriminable from them, only from those BC-verifying cases that also verify B_{disp}C can B_{disp}C be guaranteed in the wider subset of phenomenally indiscriminable cases and phenomenally-discriminable-yet-non-phenomenally-indiscriminable cases to them. This stronger guarantee has it that phenomenal discriminability does not suffice to undermine the propagation of dispositions. This is presumably because (BC ∧ B_{disp}C)-verifying cases are those in which one's disposition to believe therein is strong/stable enough to more reliably persist in other cases one can phenomenally discriminate from one's own case of actualised disposition. This can be compared to the relatively attenuated strength/stability of dispositions unactualized in cases in which BC does not obtain, and how, given D-B_{M+}, such cases can infer belief in other cases *without* even an underlying disposition. Furthermore, we can also note how there is a lesser chance of a disposition obtaining in cases phenomenally discriminable from cases in which one believes without an underlying disposition, due to DOX-B_{M+}, perhaps due to such (DOX-B_{M+})-relevant cases being those in which one is less likely to have a persistent doxastic disposition in the first place.

What would this non-phenomenal mode of discriminability be precisely? Without constructing too much *ad hoc* specification, we can tie it closely to the term of doxastic disposition and call it *dispositional* discriminability. Here, case A is dispositionally discriminable from case B iff one has a disposition to phenomenally discriminate B from A. This is different from A being phenomenally discrimin*able* from B, which implies that S *can* be liable in A to believe that C despite S not being liable in B to also believe that C. More specifically, from content relation (*m*), A being phenomenally discriminable from B entails that their contents, which matter for whether C is true or not, attain a relation such that S has the *capacity* to discern what C's truth value is in either case. Having a disposition, on the other hand, to phenomenally discriminate B from A, which amounts to a disposition to discern C's truth value in the above fashion, is not necessarily identical to one having the ability and/or capacity for such discernment, because one can have an ability to act while for whatever reason not being disposed to do so. To simplify and differentiate, phenomenal discriminability may entail a capacity that is directed to its object, that being phenomenal content, while dispositional discriminability may entail a disposition that is directed to its own object, that being *non*-phenomenal content, and where something about a case's non-phenomenal content matters for someone's disposition to phenomenally discriminate between two cases regardless of their actual capacity for such discrimination.[34]

The relation between phenomenal and dispositional discriminability is not unlike that between beliefs and dispositions then, for just like believing without an underlying disposition, one can discriminate between two cases without a corresponding disposition to do so—e.g., they can discriminate, but only luckily so, or not reliably. However, the opposite does not seem true, in that one not being able to discriminate between two cases should constitute one not having a corresponding disposition, for although a disposition to discriminate can be left unactualized, or even missing, in instances wherein one still has a capacity to discriminate, a disposition that can *never* be actualised in instances wherein one has no capacity to discriminate does not strike as an extant disposition at all. A natural interpretation of this is that a disposition to discriminate obtains only if there is a capacity to do so, meaning that dispositional discriminability entails phenomenal discriminability but not *vice versa*.

With this, let us introduce a content margin, m^*, for dispositional indiscriminability wherein A and B are dispositionally indiscriminable for S iff S cannot dispositionally discriminate—i.e., does not have a disposition to discriminate—between A and B. Being a content margin, m^* functions similarly to m, in that dispositional indiscriminability hinges on the propositional contents of cases, whereby α_i^a and α_{i+1}^b are mutually dispositionally indiscriminable for S iff $|a - b| \leq m^*$. However, they are related via $m \rightarrow m^*$. We can now specify DDOX-B_{M+}, in that, $\forall[(a, b, \mathrm{i}, \mathrm{x}, \delta) : (|a - b| \leq m^*), (\delta_i \Leftrightarrow_{N\backslash PN} \delta_x), (0 \leq \mathrm{i} \leq \mathrm{n})]$,

$$\text{DDOX-B}_{M+} \Leftrightarrow ([\delta\text{-}(BC)_i^a \wedge \delta\text{-}(B_{disp}C)_i^a] \rightarrow \delta\text{-}(B_{disp}C)_x^b).$$

[34] Referring to Note 7, one way of talking about non-phenomenal content mattering for dispositional discriminability is in externalist terms, wherein, say, such a disposition is causally connected to features of cases that bear on their non-phenomenal content—i.e., an absent disposition to phenomenally discriminate being grounded on dispositional indiscriminability as a function of agent-independent case content. A more internalist picture may have it that non-phenomenal content itself is directly linked to one's disposition to discriminate without any causal involvement with underlying case features, perhaps even to the point that the disposition being directed to its non-phenomenal object is a solely agent-indexed matter—i.e., dispositional indiscriminability about case contents being agent-dependent by being grounded on an absent disposition for the agent to phenomenally discriminate.

With DDOX-B_M+ and DD-B_M+, the anti-luminist has a valuable explanation for how beliefs and dispositions start falling off, given identical antecedents, at different discriminative and modal points. This issues in a view of the modal landscape surrounding case-series involving either α_I or δ_x that makes the case for the anti-luminist in a manner that appeals to something other than an auspicious discriminative capacity to not have R* succumb to a sorites.

VII

Now, DDOX-B_M+ concords with the three things the belief/disposition distinction must capture, given at the end of Section V: not trivializing A+, not having the case-pairs to which DD-B_M+ is irrelevant arbitrarily become those to which DOX-B_M+ and D-B_M+ are also irrelevant, and having $B_{disp}C$-verifying cases generally be closer to α_{i+2} than either BC- or (BC ∧ $B_{disp}C$)-verifying ones. Before addressing the trivialisation issue, let us see how DDOX-B_M+ permits the second and third options. First, it allows both BC and $B_{disp}C$ for some (δ_{x+2})-case that is either not sufficiently similar or not content identical to α_{i+2} yet still N\PN-related to and indiscriminable from α_{i+1}. This is because any such (δ_{x+2})-case that verifies BC can easily be N\PN-related to and dispositionally indiscriminable from α_{i+1}.[35] Next, it allows for $B_{disp}C$-verifying (δ_{x+2^+})-cases that are modally closer to α_{i+2} than those verifying some other combinations of doxastic states, because DDOX-B_M+ itself just infers $B_{disp}C$ in the relevant cases, not BC. This speaks to a greater degree of needed modification in the modal landscape surrounding $\alpha_{0\text{-}n}$, due to rules such as DDOX-B_M+, to accommodate results like α-$(\neg BC)_{i+2^c}$.

To describe what such modifications entail, let us note how the ways of relating BC-, $B_{disp}C$-, and (BC ∧ $B_{disp}C$)-verifying (δ_{x+2^+})-cases introduced by DDOX-B_M+ are reasonably explainable. Besides the already established $B_{disp}C$-verifying ones being modally closer than the BC-verifying ones, we now have a way of positionally differentiating (BC ∧ $B_{disp}C$)-verifying cases from BC-verifying ones, at least from the vantage point of α_{i+2}.Without DDOX-B_M+ one cannot establish this. To see how, first let us consider the BC-verifying (δ_{x+2^+})-case that is the consequent of D-B_M+ given α-$(B_{disp}C)_{i+2^c}$. This case is N\PN-related to α_{i+2}, and if it also does not verify $B_{disp}C$, which is possible given the relational leniency of (δ_{x+1^+}, δ_{x+2^+}) case-pairs, then, given just DOX-B_M+, any general case from which $B_{disp}C$ could otherwise be entailed for said (δ_{x+2^+})-case must be either not N\PN-related to or not phenomenally indiscriminable from it. This would obtain regardless of the general case being BC-

[35] DDOX-B_M+ does not necessarily prevent (α_{i+1}, δ_{x+2^+}) case-pairs from expressing mere disposition-loss, even if they are phenomenally indiscriminable, because not all such pairs must be N\PN-related.

or (BC ∧ B_{disp}C)-verifying. However, with DDOX-B_{M+} in the picture, any such (BC ∧ B_{disp}C)-verifying general case N\PN-related to said (δ_{x+2}^{+})-case would have to be dispositionally discriminable from it, while any (BC ∧ ¬B_{disp}C)-verifying one N\PN-related to it would only need to be phenomenally discriminable from it. If this BC-verifying (δ_{x+2}^{+})-case were instead also B_{disp}C-verifying, then the order switches: (BC ∧ B_{disp}C)-verifying cases come first starting among cases content-identical to it, while (BC ∧ ¬B_{disp}C)-verifying ones start being found among cases dispositionally discriminable from it.

However, not all (δ_{x+2}^{+})-cases N\PN-related to α_{i+2} are those within which D-B_{M+} entails BC. Some can be (¬BC ∧ ¬B_{disp}C)-verifying while others are (¬BC ∧ B_{disp}C)-verifying. Here, depending on which one, the order between (BC ∧ ¬B_{disp}C)- and (BC ∧ B_{disp}C)-verifying cases related to any one of these (δ_{x+2}^{+})-cases can change. If δ_{x+2}^{+} is (¬BC ∧ B_{disp}C)-verifying, then, due to DD-B_{M+}, among cases sufficiently similar to it, (BC ∧ ¬B_{disp}C)-verifying cases come first starting among cases content-identical to δ_{x+2}^{+}, while (BC ∧ B_{disp}C)-verifying ones start being found among cases content-non-identical to it. If δ_{x+2}^{+} is (¬BC ∧ ¬B_{disp}C)-verifying, then the order is the same as when δ_{x+2}^{+} is (BC ∧ ¬B_{disp}C)-verifying.

How this relates to α_{i+2} hinges on how likely δ_{x+2}^{+} is to verify a particular doxastic state. Given the positional closeness between δ_{x+1}^{+} and δ_{x+2}^{+}, one would expect (δ_{x+1}^{+}, δ_{x+2}^{+}) being sufficiently similar to be more common than it being not N\PN-related, thereby rendering δ^{+}-$(B_{disp}C)_{x+2}{}^{c}$ more likely true than false. Indeed, given that B_{disp}C-verifying cases are modally closer to α_{i+2} than BC-verifying ones, one could expect (¬BC ∧ B_{disp}C)-verifying (δ_{x+2}^{+})-cases to be closer than (BC ∧ B_{disp}C)-verifying ones to α_{i+2}. From this we can determine that, while considering that BC-verifying (δ_{x+2}^{+})-cases cannot be sufficiently similar to α_{i+2}, among the cases sufficiently similar to α_{i+2}, there ought to be more (BC ∧ ¬B_{disp}C)-verifying (δ_{x+2})-cases phenomenally indiscriminable from it than (BC ∧ B_{disp}C)-verifying ones.

Interestingly enough, from the vantage point of α_{i+1}, the relationship changes. Because we are not required to guarantee something like α-$(\neg BC)_{i+1}{}^{b}$, given DOX-B_{M+} and DD-B_{M+} the closest kinds of general cases to α_{i+1} would be (BC ∧ B_{disp}C)-verifying ones guaranteed starting among cases content-identical and sufficiently similar to α_{i+1}. This would be followed, given DOX-B_{M+}, by just B_{disp}C-verifying ones starting among cases content-non-identical and/or not sufficiently similar to α_{i+1}, then finally, given DDOX-B_{M+}, by just BC-verifying ones starting among all cases non-N\PN-related and/or dispositionally discriminable to α_{i+1}.

The main take away here is that, by respecting the fact of dispositions and beliefs falling off at different points, the type of cases – (BC ∧ B_{disp}C)-verifying ones – closest to the $\alpha_{0\text{-}n}$ case-series at one portion must be further from it at some other portion that is needed to halt a sorites.[36] This is contentious given that, from the vantage point of having a disposition to believe, intuitively, instances of belief without an underlying disposition should be less strongly related to the actual case of disposition, whether α_{i+1} or α_{i+2}, than instances of a belief resulting from an underlying disposition, for dispositions are not even found in the former type of instances at all.

Now, the luminist can avoid this issue altogether. They can equate BC-verifying cases with ($BC \wedge B_{disp}C$)-verifying ones and require that, at the point in the actual series where ($\neg BC \wedge \neg C$) must obtain, all cases attaining a particular relation (modal/discriminative) to such a point simply verify ($\neg BC \wedge \neg B_{disp}C$) as well. In other words, this amounts to making beliefs and dispositions rise and fall together, or conflating phenomenal and dispositional indiscriminability, at least to the extent needed to safeguard $\alpha\text{-}(\neg BC)_{i+2}{}^{c}$. By this, the luminist can prevent the contentious modal modification whereby relatively more intuitively plausible doxastic states – e.g., ($BC \wedge B_{disp}C$) – become closer in modal space to relatively less intuitive ones – e.g., ($BC \wedge \neg B_{disp}C$). This response is even better suited as a luminist response than an anti-luminist one, for the former can leverage one's auspicious discriminative capacities to explain it while the latter must jerrymander the modal landscape to attain it. Moreover, this can happen even without the luminist having to deny A+ for of this explanation.

However, note the nuance of the anti-luminist's position here. They expressly deny the luminist's move of appealing to an auspicious discriminative capacity because it is not this that, for the anti-luminist, avoids sorites. If one's discriminative capacity alone cannot do the job, then obviously the modal landscape must play its part by manifesting in the required manner. The fact that the above combination of doxastic states that must happen to be sensitive to the belief/disposition distinction is contentious actually works in the anti-luminist's favour, as it all the more serves as reminder that results like $\delta^{+}\text{-}(BC \wedge \neg B_{disp}C)_{i+2}{}^{c}$ are not as possible as $\delta^{+}\text{-}(\neg BC \wedge B_{disp}C)_{i+2}{}^{c}$, which goes against luminist intuitions about some easy equivalency between phenomenal and dispositional indiscriminability. Indeed, all the anti-luminist needs to guarantee is for this *change* in the doxastic state of affairs—e.g., that ($BC \wedge \neg B_{disp}C$) obtains for general cases that *become* closer to other cases verifying ($BC \wedge B_{disp}C$) —to happen at least once in the case-series, while the luminist must guarantee a discriminative capacity that can apply to any point in the case-series at which BC would otherwise be unsafe—i.e., the anti-luminist just needs one

[36] This takes place even if $\alpha\text{-}(\neg BC)_{i+1}{}^{b}$ is true, because then the modal/discriminative requirements simply move back a position in the case-series – i.e., those for α_{i+2} and α_{i+1} when $\alpha\text{-}(BC)_{i+1}{}^{b}$ is true become those for α_{i+1} and α_{i} when $\alpha\text{-}(\neg BC)_{i+1}{}^{b}$ is true, respectively.

such instance and is not forced to give an explanation for why it could happen more times, but the luminist must employ one's discriminative capacities in a way that must also be able to explain multiple such instances. The luminist's guarantee obviously requires more than the anti-luminist's and thus strengthens the latter's position.

With this we can finally move on to discussing trivialising safety. The way the anti-luminist can conceive of not trivialising A+ is through DDOX-B_M+, specifically m^*, and the possibility of δ_{x+2^+} and α_{i+2} being phenomenally discriminable yet dispositionally indiscriminable from α_{i+1}. This is because, by this possibility, the anti-luminist can accept the luminist intuition of (α-(BC)$_{i+1}{}^b \wedge \alpha$-($\neg$BC)$_{i+2}{}^c$) obtaining due to $|b-c| > m$ without conceding that (α-(KC)$_{i+1}{}^b \wedge \alpha$-($\neg$C)$_{i+2}{}^c$) obtains also due to $|b-c| > m$.[37] In other words, the anti-luminist can avoid begging the question against luminist intuitions about phenomenal discriminability without denying margins, because even if α_{i+2} is phenomenally discriminable from α_{i+1}, it may be dispositionally indiscriminable from α_{i+1}. This means that the resulting relatively unstable disposition, or lack thereof, to discriminate α_{i+2} from α_{i+1} could entail the existence of ($\neg$C $\wedge$ BC)-verifying cases that are N\PN-related to and phenomenally indiscriminable from α_{i+1}, much to the luminist's ire—i.e., you being able to discriminate α_{i+2} from α_{i+1} in terms of the truth of C does not ensure that you always will discriminate in this fashion and have α-(BC)$_{i+1}{}^b$ be perfectly safe, because not having a reliable disposition to discriminate allows for a persistent disposition to believe that C, thereby motivating inclusion within the set of (A+)-relevant BC-verifying cases also some of those that verify $\neg$C.[38] Additionally, cases that are N\PN-related to, and phenomenally *discriminable* from α_{i+1} may also be ($\neg$C $\wedge$ BC)-verifying if they are still dispositionally indiscriminable from α_{i+1}, which paints as more suspicious the luminist intuition of (α-(BC)$_{i+1}{}^b \wedge \alpha$-($\neg$BC)$_{i+2}{}^c$) obtaining due to phenomenal discriminability between the relevant cases.[39]

---

[37] Now, this still ensures α-(B_{disp}C)$_{i+2}{}^c$, but through DDOX-B_M+, not DOX-B_M+, for (α_{i+1}, δ_{x+2^+}) being a dispositionally indiscriminable pair entails (α_{i+1}, α_{i+2}) being one as well given content identity between α_{i+2} and δ_{x+2^+}. Note also that, even with DDOX-B_M+, (α_{i+2}, δ_{x+2^+}) must remain an insufficiently similar case-pair for any BC-verifying (δ_{x+2^+})-case even if (α_{i+1}, δ_{x+2^+}) turns phenomenally discriminable albeit dispositionally indiscriminable, for otherwise, by DDOX-B_M+, any such BC-verifying (δ_{x+2^+})-case would become B_{disp}C-verifying and contravene the initial assumption of α-($\neg$BC)$_{i+2}{}^c$.

[38] This result would persist without requiring for safety the absence of false beliefs in all N\PN-related and *dispositionally* indiscriminable cases. Additionally, such a requirement is unreasonable, for dispositional discriminability is being deployed to mark the threshold where doxastic *dispositions* start falling off, while A+ is written in terms of *belief*, not dispositions.

[39] This pushes back against the luminist rejoinder that safe beliefs require some capacity to discern,

The luminist can reject this assessment, of course, but only at a heavy cost. First, although DDOX-B_M+ is not needed to derive R*, it is perhaps just as plausible as DOX-B_M+, so the luminist rejecting it would be committed to also reject DOX-B_M+, an already intuitively plausible principle in its own right. Second, if the luminist instead asserts that no (A+)-relevant BC-verifying case is also (¬C)-verifying, then they must commit to either one of three problematic points. One, they conflate beliefs and dispositions to believe by requiring a dividing line, within the set of cases N\PN-related to α_i, between (BC ∧ B_{disp}C ∧ C)-verifying cases phenomenally indiscriminable from, or even content-identical to α_i and (¬BC ∧ ¬B_{disp}C ∧ ¬C)-verifying ones phenomenally discriminable from α_i. Two, they make dispositions to believe have nothing to do with belief, such that a (BC ∧ C)/(¬BC ∧ ¬C) dividing line becomes independent from a B_{disp}C/¬B_{disp}C dividing line. In other words, these two points amount to the *ad hoc* suppositions that, one, $m = m^*$, or even $m = m^* = 0$, is true, or two, $(m \rightarrow m^*)$ is false. Three, by instead accepting DDOX-B_M+, $m \neq m^*$, and $(m \rightarrow m^*)$, the luminist would be forced to resist R* by simply supposing that, for purely modal reasons, (A+)-relevant BC-verifying cases are always C-verifying—i.e., any (¬C ∧ BC)-verifying case phenomenally indiscriminable from α_{i+1} is never N\PN-related to it.

The problem with the first cost should be obvious if the luminist is thereby forced to reject DOX-B_M+.[40] Points one and two of the second cost trivially make both B_M+ and A+ true for any N\PN-related (α_{i+1}, δ_{x+2}^+) case-pair, thus having α-$(BC)_{i+1}{}^b$ always be safe through perfect discriminative capacities such that content differences between cases within any N\PN-related (α_{i+1}, δ_{x+2}^+) case-pair both are greater than m and also either are greater than m^* or have no bearing on m^* whatsoever. Both ways of relating m and m^* are highly contentious since they equate

as Vanrie (2020) puts it, that things "are different in sufficiently similar cases," but not that things are different in "sufficiently similar ([phenomenally] indiscriminable) cases" (549), for the possibility of (¬C ∧ BC)-verifying cases that are N\PN-related to and phenomenally discriminable from the actual present case of belief can speak against one's capacity to discriminate between these phenomenally discriminable cases as sufficient for beliefs being safe.

[40] Not just with DOX-B_M+, but rejecting any one of the principles that jointly establish R* would be worrying: D-B_M+ has very minimal requirements on the modal space; DOX-B_M+ is intuitively plausible on account of its falsification implying instances of belief without disposition in at least a subset of close by possible cases to an actual belief-case, which seems implausible enough especially given the modal analysis done in this Section; DD-B_M+ and DDOX-B_M+ work in tandem to distinguish beliefs from dispositions, so denying any one of the two risks, as has already been discussed, contentious conflations; DUPL-N/I-3 is a simple transitivity relation that prevents sorites from surfacing; and A+ is a somewhat general way of formalising the link between knowledge and belief in actual *and* possible cases, so it is pretty reasonable as safety principles go.

to either a mutual relation of conflation or one of irrelevance. Indeed, the point of having beliefs start falling off sooner than dispositions to believe is so that dividing lines like (BC ∧ C)/(¬BC ∧ ¬C) become highly unmotivated due to an incursion of B_{disp}C-cases into the set of (¬BC ∧ ¬C)-cases—an incursion of greater extent than what would obtain if beliefs and dispositions start falling off at the same time, or if dispositions start falling off earlier than beliefs. Lastly, point three of the second cost is worse off than the other two since the other two still allow for beliefs and dispositions to fall off for discriminative reasons first before modal reasons—the first two points at least respect the fact, argued for in the prior Section, that it is more ostensibly plausible for an N\PN-related case to also be discriminable from α_i than for an indiscriminable case to not be N\PN-related to it. On the other hand, point three entails beliefs and dispositions falling off for modal reasons first before discriminative reasons, trivially making both B_M+ and A+ true for any phenomenally indiscriminable (α_{i+1}, δ_{x+2}) case-pair and thus having α-$(BC)_{i+1}{}^{b}$ always be safe by dint of an auspicious modal landscape. Obviously, these costs are problematic and motivate the deployments of the anti-luminist instead.

In conclusion, the anti-luminist prevents a sorites problem concerning beliefs and doxastic dispositions being implicated from R* by motivating a modal/discriminative analysis about beliefs and dispositions falling off at different places. Crucially, this result obtains even without perfect discriminative capacities, phenomenal or dispositional. The luminist seeking to falsify R* must therefore stop a sorites by appealing either to a discriminative capacity that ensures results like (α-$(BC)_{i+1}{}^{b}$ ∧ α-$(\neg BC)_{i+2}{}^{c}$) or an auspicious modal landscape that does the same, all while arguing for α-$(KC)_{i+1}{}^{b}$. As discussed above, how the luminist gets about accomplishing this involves worrying ramifications, concerning the relation between beliefs and their underlying dispositions, of overly strong appeals to either just our discriminative capacities or just an auspicious modal landscape. Alternatively, the anti-luminist settles for a more conservative approach by leveraging the influence of both in having beliefs fall off sooner than dispositions for either modal or discriminative reasons, which explains how beliefs like α-$(BC)_{i+1}{}^{b}$ can be rendered unsafe due to the persistence of B_{disp}C-verifying cases causing some (BC ∧ ¬C)-cases to be (A+)-relevant. This is because, otherwise, dividing lines like (BC ∧ C)/(¬BC ∧ ¬C) can be motivated by citing the resultantly relatively smaller influence of B_{disp}C-cases within the set of (¬BC ∧ ¬C)-cases, thereby grounding the luminist's desired possibility of (α-$(KC)_{i+1}{}^{b}$ ∧ α-$(\neg C)_{i+2}{}^{c}$) on the (BC ∧ C)/(¬BC ∧ ¬C) division representing an auspicious enough discriminative capacity ensuring doxastic safety.

VIII

Now, this does not mean that there are no possible luminist rejoinders to margins principles that are applicable as counters to my arguments above. Indeed, many of these responses can be interpreted as attempts at rationalising conflation of belief and disposition, at least in ways whereby the fact of them falling off at different places no longer is relevant to the safety principles at play. How might a luminist motivate this manoeuvre? One method of doing so that is discussed below deals with a variant of the rule, BC → C, meant to safeguard results like $\alpha\text{-}(KC)_{i+1}{}^{b}$ due to $\alpha\text{-}(BC)_{i+1}{}^{b}$ being an instance of safe belief. This manoeuvre differs from the previous attempts by the luminist at denying R* and trivially verifying A+, because the manoeuvre's non-trivial accommodation of A+ can be grounded on appeals to other safety principles besides A+. Obviously, (BC → C) being true would be disastrous for the anti-luminist, because our modal analysis above hinges on a belief/disposition distinction requiring mere belief-loss in $(\alpha_{i+1}, \alpha_{i+2})$ being distinct from possible mere disposition-loss in $(\delta_{x+1}{}^{+}, \delta_{x+2}{}^{+})$, which is denied if BC → C is true given that $\delta^{+}\text{-}(BC \wedge \neg C)_{x+2}{}^{c}$ would then be impossible. Therefore, (BC → C) being true implicates R* as question-begging against the luminist. As such, for my analysis to resist this, any examination of luminist espousals of BC → C must therefore indicate how they are misguided.

Berker's use thereof is informative. BC → C is employed by Berker (2008) to describe all situations that are also applicable to the luminosity principle, L* ⇔ (C → KC), (18) in that it comprehends those cases of BC that are *knowledgeable*, even those the anti-luminist would not consider so. Specifically, Berker constraints BC → C through his rule, **L**$_C$, wherein one's BC is factive "[i]f one has done everything one can to decide whether" C obtains.(9) If we assume that the antecedent obtains for every situation in which L_C applies, this being reminiscent of the transition in Section I from talk about 'position to know' to that of 'knowledge', (Wong 2008, 537) then we can rewrite L_C such that, $\forall[(a, i, \delta) : (0 \leq i \leq n)]$,

$$L_C{}^{*} \Leftrightarrow [\delta\text{-}(BC)_i{}^{a} \rightarrow \delta\text{-}C_i{}^{a}].$$

Now, L_C*, although it being tied to luminosity in general entails it having to be inconsistent with R*, still can be further interpreted in either a weak or strong way, having to do with how universally applicable the rule is. For instance, a strong interpretation (call it S-L_C*) would entail that (BC ∧ ¬C)is simply false for any general case, while a weaker variant (call it W-L_C*) would allow some instances of (BC ∧ ¬C) while only requiring, say, that $\alpha\text{-}(BC)_{i+1}{}^{b}$ counts as knowledgeable belief. S-L_C* is obviously inconsistent with the doxastic rationale behind R* because it contradicts

B_M+ for any (α_{i+1}, δ_{x+2}^{+}) case-pair by essentially making safety, A+, trivial.[41] Nonetheless, although W-L_C+ is consistent with B_M+ for any (α_{i+1}, δ_{x+2}^{+}) case-pair, it contradicts R* as it goes against A+ by identifying unsafe true belief as knowledgeable in some instances, depending on other safety principles at play.

W-L_C* is important to avoid begging the question against the anti-luminist by stipulating perfect doxastic truth-reliability. To this end, Berker (2008) proposes the obtaining of "a tight connection between the obtaining of certain conditions and our beliefs, at least upon reflection, about the obtaining of those conditions." (17) This connection, by assuming again the above transition motivating L_C*, speaks to W-L_C* rendering at least some borderline cases of true belief—e.g., α-$(BC)_{i+1}{}^{b}$—more appropriate for knowledge than not, with the extent of such cases linked to the strength of the connection's tightness. In other words, depending on said strength, W-L_C* could permit something like δ^{+}-$(BC \land \neg C)_{x+2}{}^{c}$ without it sufficing as falsification of α-$(KC)_{i+1}{}^{b}$ due to α-$(BC)_{i+1}{}^{b}$ being unsafe in terms of A+. This could be permitted even with (α_{i+1}, δ_{x+2}^{+}) being an N\PN-related case-pair as long as either safety is rejected altogether for knowledge or simply another safety principle besides A+ is employed instead. Since the former option risks trivialising knowledge for *all* borderline cases of true belief, the latter option becomes the most plausible. This latter option also looks to be what Barz (2017) is agreeing with when noting that the right conclusion of accepting a margins principle at the expense of luminosity is not "that it is possible for someone who feels cold to introspect as assiduously as possible without thereby coming to *know* that one feels cold, but rather… that it is possible for someone who feels cold to introspect as assiduously as possible without thereby coming to *safely believe* that one feels cold." (482) Here, while the luminist could accept that α-$(BC)_{i+1}{}^{b}$ is unsafe in terms of A+, the anti-luminist would not be able to also conclude that α-$(KC)_{i+1}{}^{b}$ fails without presupposing A+ to be the safety principle at play.

However, and this is the crux of the issue, this appeal to W-L_C* would not work for the luminist. The anti-luminist could simply repurpose the above account of the belief/disposition distinction to accommodate some other safety principle besides A+ and reinterpret R* as needed. In fact, this strategy would even be useful to establish margins principles of a wider scope than just R*. The general strategy is as follows: generalise the modal/discriminative relations involved in the rules that establish R* while also being sensitive to beliefs and dispositions falling off at different thresholds. Now, DD-B_M+, D-B_M+, DOX-B_M+, DUPL-N/I-3, and A+ jointly

[41] This is one way of interpreting how luminosity-adjacent principles can resist a margins requirement for knowledge without necessarily resisting a safety requirement as well. See Stalnaker (2019, 32-34).

establish R*, given conditions outlined in Section IV, while DDOX-B_M+ entails that R* does not beg the question against luminist intuitions of phenomenal discriminability. With this, let us first consider three relations (r1, r2, r3) each composed of modal and discriminative dimensions, r1* and $r1^+$, and r2* and $r2^+$, etc., respectively. These are related as $r2^*/r2^+ \leq r3^*/r3^+ \leq r1^*/r1^+$– e.g., N\PN-relation as r1*, phenomenal indiscriminability as $r1^+$, sufficient similarity as r2*, and content identity as $r2^+$. Let us also introduce another discriminative relation, $r0^+$, whose strength is related in the fashion of $r1^+ \leq r0^+$. Here, (δ_x)- and (δ_x^+)-cases would be r1*- and ($r2^+ \wedge r1^*$)-related to α_i, respectively. Let us then construct five general principles: first, BC infers $B_{disp}C$ in those cases exhibiting an r1-relation to the BC-case; second, $B_{disp}C$ infers BC in at least one case ($r2^+ \wedge r1$)-related to said $B_{disp}C$-case; third, (BC ∧ $B_{disp}C$) infers either BC or $B_{disp}C$ in cases r2-related or ($r1^* \wedge r0^+$)-related, respectively, to the (BC ∧ $B_{disp}C$)-case; fourth, r1* is fully transitive over r2*, and $r1^+$is fully transitive over, in parts, $r1^+$ and $r2^+$only if $r2^+$expresses content-identity; lastly, fifth, KC requires (BC → C) in cases r3-related to the KC-case. These five principles jointly entail that KC is factive over those cases r3-related to the KC-case that are also r2-related to any BC-case either r1-related to the KC-case or ($r1^* \wedge r0^+$)-related to a (BC ∧ $B_{disp}C$)-case. These principles are also generalisations of the rules used to establish R*, so their plausibility at least has some initial backing—e.g., r1* being fully transitive over r2* functions analogously to DUPL-N/I-3 in preventing sorites.

Now, the luminist is incentivised to choose some relational content for r3 such that cases verifying (BC ∧ ¬C) are necessarily (¬r3)-related to α_{i+1}, because such a choice contradicts R* by safeguarding KC, via W-L_C*,for cases r3-related to α_{i+1}. This would also make good on a plausible luminist intuition that, as Neta and Rohrbough (2004) indicate, some general (BC ∧ ¬C)-cases, "which are initially similar in just about every respect" to cases such as α_{i+1} "except for the truth of the proposition believed," should "not prevent [one] from having knowledge in [α_{i+1}]", at least insofar as similarity is not being measured "in terms of the truth of the proposition believed." (399, 404)

However, the anti-luminist attains a dialectical advantage here. Recall the relational analysis of BC-, $B_{disp}C$-, and (BC ∧ $B_{disp}C$)-verifying (δ_{x+2}^+)-cases presented in the previous Section. An analogous analysis applicable here would be that BC-verifying general cases start coming up in cases r2- to (¬r1)-related to α_{i+1}. This means that the luminist's move forces the r3/¬r3 divide to function as a C/¬C divide within the set of BC-verifying cases. However, if the luminist accepts that these divides do not also represent a corresponding $B_{disp}C$/¬$B_{disp}C$ divide,[42] then, given that

[42] If the luminist does not accept this, then not only would they be making problematic conflations over a number of variables, but they would also be assuming that r1 ≡ r3 given how BC entails

beliefs fall off sooner than dispositions to believe and that luminists do not beg the question against anti-luminists by denying this, we can expect the luminist to concede that a few cases within the (¬C)-side of the C/¬C divide also verify (BC ∧ B_{disp}C). Depending on the extent of such (B_{disp}C ∧ ¬C)-verifying cases in the (¬r3)-side, the anti-luminist could thereby argue that the r3 relation does not satisfactorily capture a viable safety principle for belief if the extent of such cases is enough to where, intuitively, one would think that these cases ought to matter to the safety of α-$(BC)_{i+1}{}^{b}$—i.e., the extent of such (B_{disp}C ∧ ¬C)-verifying cases is large enough to where it would be more plausible for a (BC ∧ ¬C)-verifying case to be r3-related to α_{i+1} than not, for otherwise one would need to explain why the only place in which you find this large of a presence of a disposition to falsely believe that C is conveniently in safety-irrelevant cases.

Note the manoeuvre here. If a doxastic safety principle applies over the r3-relation, but the luminist does not admit that one's discriminative capacities suffice for equating the r3/(¬r3) divide with the (BC ∧ C)/(¬BC ∧ ¬C) divide—i.e., sufficing for trivial verification of the safety principle by making (BC ∧ ¬C) impossible on either side of the divide—then depending on the extent of (BC ∧ ¬C)-cases within the ¬r3 side, currently conceptualised, one could have a stronger or weaker argument for reconceptualising the r3-relation to include such cases and deem a particular borderline case of true belief unsafe, such as α-$(BC)_{i+1}{}^{b}$. The anti-luminist thus attains a dialectical advantage over the luminist, not because of some faculty of precisely identifying the extent of said (BC ∧ ¬C)-cases, but merely because this extent represents a non-null set of cases that also includes B_{disp}C-verifying ones. This is because, even without a direct way of ascertaining the extent of (BC ∧ ¬C)-cases, it is known, through D-B_{M+} and DD-B_{M+}, that this extent is already larger, due to the inclusion of B_{disp}C-verifying cases, than what would have otherwise obtained had B_{disp}C-verifying cases not been included. To claim that there is no motivation for reconceptualising the r3-relation by virtue of the presence of B_{disp}C-cases is to claim that a belief is always safe despite any extent to which one could be said to have a disposition to believe falsely in nearby cases, which just sounds problematically strong as a claim. Now, if this reconceptualised r3 is called r4, such that r3 ≤ r4 ≤ r1, the luminist can still respond to the anti-luminist by having cases that verify (BC ∧ ¬C) also be necessarily (¬r4)-related to α_{i+1}. However, this would just motivate the above manoeuvre again, favouring a reconceptualization of the r4-relation and thus the anti-luminist position of having α-$(BC)_{i+1}{}^{b}$ be unsafe.

Indeed, the luminist does not have many options here. Given that α-$(BC \wedge B_{disp}C)_{i+1}{}^{b}$ infers B_{disp}C in all cases (r1* ∧ r0⁺)-related to α_{i+1}, even if the luminist

---

B_{disp}C over all cases r1-related to the BC-case.

appeals to some r5 relation for safety such that, one, $r1^+ \leq r5^+ \leq r0^+$, and two, cases verifying (BC ∧ ¬C) are necessarily (¬r5)-related to α_{i+1}, there will still be cases verifying ($B_{disp}C$ ∧ ¬C) on the ¬r5 side that could restart the entire above argument once again to have r5 reconceptualised and thus $\alpha\text{-}(BC)_{i+1}{}^{b}$ remain unsafe. The only option, it seems, available to the luminist would be to appeal to some further relation, r6, such that $r0^+ \leq r6^+$, have KC require the absence of false BC in all cases r6-related to the KC-case, and have cases verifying (BC ∧ ¬C) be necessarily (¬r6)-related to α_{i+1}. By this move, the extent of $B_{disp}C$-cases in the ¬r6 side could be reduced to the point where the luminist would have a better argument to not count the (BC ∧ ¬C)-cases as mattering to doxastic safety by being r6-related to α_{i+1}; after all, $\neg r0^+$ instantiates sooner than ¬r6 and $B_{disp}C$ starts falling off at cases ($\neg r0^+$)-related to the (BC ∧ $B_{disp}C$)-case. As a result, this safeguards W-L_C* for all cases relevant to the safety of $\alpha\text{-}(BC)_{i+1}{}^{b}$ and thus to the grounds for $\alpha\text{-}(KC)_{i+1}{}^{b}$ and, consequently, ¬R*. However, even if the anti-luminist can no longer leverage the set of $B_{disp}C$-verifying cases in the ¬r6 side being extensive enough to count the (BC ∧ ¬C)-verifying cases as safety-relevant once again, the fact that the luminist must at the very least equate the C/¬C divide within the set of BC-verifying cases with the divide marking the threshold at which $B_{disp}C$ starts falling off is still contentious. Given that beliefs fall off sooner than dispositions, (BC ∧ ¬C)-verifying cases at the ¬r6 side start having less to do with a strong disposition to believe in a false proposition and more to do with false beliefs being a far-enough possibility that they cannot undermine the safety of $\alpha\text{-}(BC)_{i+1}{}^{b}$, probably because ¬C itself is a far-enough possibility in relation to α_{i+1}, albeit not as far as (¬BC ∧ ¬C). Nevertheless, this is a weak argument by the luminist, especially since they require it to guarantee the safety of BC in α_{i+1}, because α_{i+1} neighbours the (¬C)-verifying case of α_{i+2}, which does not seem at all like a *far-enough* possibility for α_{i+1}. As such, the anti-luminist's establishment of R* should still stand.

The dialectical advantage conferred to the anti-luminist is therefore based on the luminist accepting that, one, safety cannot be trivially guaranteed by one's discriminative capacities, and two, beliefs start falling off sooner than their underlying dispositions. With these acceptances, the luminist is thereby forced to allow for the possibility that the set of nearby cases of false belief is extensive enough, due to these cases being informed in part by an underlying disposition to have false beliefs, to once again be relevant to discussions surrounding the safety of borderline cases of true belief—e.g., $\alpha\text{-}(BC)_{i+1}{}^{b}$. Establishing that such borderline cases are indeed safe thus necessitates extending the safety principle's modal/discriminative relation to a point where it starts becoming inconsistent with the modal set-up of the ALA. The anti-luminist consequently comes out on top,[43] as the only option left for the luminist would be to espouse S-L_C* and either deny a safety requirement for knowledge or trivialise satisfaction of safety.[44] Again, this is not to say that the luminist is vying for anything new to the anti-luminist's detriment, just that what the former can argue for outside of perfect phenomenal discriminability or trivial safety can, by the analysis provided here, be accommodated by the latter in a way that favours margins for knowledge.

Furthermore, these anti-luminist responses work to justify much stronger margins claims than R*. For instance, the exact character of any of the r1-r6 relations can concomitantly change depending on the quality of our discriminative capacities—i.e., weaker capacities permitting propagation of beliefs and/or dispositions to believe over larger sets of cases. Therefore, if r3 is composed of phenomenal indiscriminability and N\PN-relation, then given the assumption that r2 ≤ r3 in terms of scope, knowledge being factive over those cases r3-related to the KC-case that are also r2-related to at least one (BC ∧ B_{disp}C)-case expresses a margins rule that is closer to R+: knowledge is factive over a sub-set of cases phenomenally indiscriminable from and N\PN-related to the KC-case.

If the luminist targets this rule establishment by having all (BC ∧ B_{disp}C)-cases r2-related to any case r3-related to the KC-case become (¬r2)-related instead, thereby allowing for some (¬C)-cases r3-related to the KC-case to falsify the broader margins rule, the anti-luminist can always leverage the above manoeuvre that renders α-$(BC)_{i+1}{}^{b}$ unsafe to motivate a reconceptualization of the r2 relation to reinstate margins once again. This would force the luminist to keep extending (BC ∧ B_{disp}C)-cases further and further away from the case-series until a point where

[43]This manner of incorporating safety into a response against the luminist may even be employed to undermine arguments in favour of the KK-thesis that conceive of the modal landscape in particular KK-sufficient ways. See Greco (2014), Das & Salow (2018), and Goodman & Salow (2018).

[44] Balog (2012) and Barz (2017) indirectly talk about S-L_C* in their discussions of mental quotation and direct phenomenal concepts, respectively. Of course, there is the possibility of what Chalmers (2003) terms "standing phenomenal concepts", (239) which would preclude factivity for belief, but the main point of principles like S-L_C* is that they are valid *for* the right types of mental mechanisms underlying one's belief-forming processes. Indeed, the ALA can be read as a way to delegitimise the margins-relevancy of S-L_C* because nearby cases may include those of false belief if one cannot distinguish between cases involving, say, direct phenomenal concepts and those occasioning mere standing phenomenal concepts. In other words, even if the luminist contends that, by an account of constitutive connections between facts of the matter and beliefs thereof, it would be unreasonable for particular *types* of beliefs to be false, to then say that cases of belief that reasonably *can* be false are necessarily outside the modal relation conditioning the safety of beliefs that *cannot* be reasonably false is to conflate perfect discriminability with trivial safety.

knowledge itself seems unwarranted. This is because, eventually, the only doxastically relevant cases close to the KC-case would be those either of unactualized dispositions or of beliefs without an underlying disposition, at which point the status, in the supposedly KC-verifying case, of BC *as knowledgeable belief* starts turning tenuous. After all, if the only close by and closely similar cases to the present case are those of dispositionally unsupported beliefs or unactualized dispositions to believe, then there is an argument to be had that whatever environment surrounds the present belief-case is inauspicious enough to reasonably prevent it from being a case of knowledge.

IX. Conclusion

An effective coarse-grained analysis of the ALA favours the anti-luminist if they can establish a margins principle by achieving two things. One, an account of belief propagation that, by being sensitive to when the propagation halts, is also not susceptible to sorites, based on rules that, together, neither beg the question against nor concede too much to the luminist. Two, an account that, when resisted by the luminist, corners them into contentious positions. My account achieves this through a distinction between beliefs and dispositions to believe that grounds the possibility of halting belief propagation on modal/discriminative requirements that are less demanding than more luminist solutions for preventing sorites. These requirements, centred on the notion that beliefs fall off sooner/easier than dispositions to believe, do away with ways of stopping sorites that appeal to discriminative capacities that trivialise the safety of beliefs. This forces the luminist to accept a relation between beliefs and dispositions through which it becomes more reasonable than not to regard cases of false belief as near-enough to actual cases of borderline belief to where the former cases affect the safety of beliefs verified in the latter cases. By doing so, the luminist must choose between conceding the truth of margins, revoking safety altogether for knowledge, or arguing for alternative safety principles. The first two choices are problematic because the luminist must either revoke luminism or accept that knowledge can be gained very easily, even in highly inauspicious environments. If the luminist chooses the third option, then the anti-luminist can re-apply their desired belief/disposition distinction to argue for the insufficiency of the luminist's newly chosen safety principle in guaranteeing the safety-irrelevance of false belief-cases for borderline belief-cases. By repeating this cycle of choices, the luminist enters into a dialectic with the anti-luminist that compels the former to side with a possible doxastic state of affairs that goes against the very spirit of the ALA's set-up. The anti-luminist consequently comes out on top with the advantage.

References

Balog, K. 2012. "Acquaintance and the Mind-Body Problem." In *New Perspectives on Type Identity: The Mental and the Physical*, edited by S. Gozzano and C. S. Hill, 16-42. Cambridge: Cambridge University Press.

Barz, W. 2017. "Luminosity Guaranteed." *Pacific Philosophical Quarterly* 98(S1): 480-496.

Berker, S. 2008. "Luminosity Regained." *Philosophers' Imprint* 8(2): 1-22.

Bonnay, D., & Égré, P. 2008. "Margins for Error in Context." In *Relative Truth*, edited by M. García-Carpintero and M. Kölbel, 103-128. Oxford: Oxford University Press.

Brueckner, A., & Fiocco, M. O. 2002. "Williamson's Anti-Luminosity Argument." *Philosophical Studies* 110 (3): 285-293.

Chalmers, D. 2003. "The Content and Epistemology of Phenomenal Belief." In *Consciousness: New Philosophical Perspectives*, edited by Q. Smith and A. Jokic, 220-272. Oxford: Oxford University Press.

Das, N., & Salow, B. 2018. "Transparency and the KK Principle." *Noûs* 52(1): 3-23.

Dokic, J., & Égré, P. 2009. "Margin for error and the transparency of knowledge." *Synthese* 166(1): 1-20.

Dorst, K. 2019. "Abominable KK Failures." *Mind* 128(512): 1227-1259.

Goodman, J., & Salow, B. 2018. "Taking a chance on KK." *Philosophical Studies* 175(1): 183-196.

Greco, D. 2014. "Could KK be OK?" *The Journal of Philosophy* 111(4): 169-197.

Immerman, D. 2018. "Kumārila and Knows-Knows." *Philosophy East and West* 68(2): 408-422.

———. 2020. "Williamson, closure, and KK." *Synthese* 197(8): 3349-3373.

Jenkins, D. 2021. "Luminosity in the stream of consciousness." *Synthese* 198(Suppl 7): S1549–S1562.

McHugh, C. 2010. "Self-knowledge and the KK principle." *Synthese* 173(3): 231-257.

Neta, R., & Rohrbaugh, G. 2004. "Luminosity and the Safety of Knowledge." *Pacific Philosophical Quarterly* 85(4): 396-406.

Ramachandran, M. 2012. "The KK-Principle, Margins for Error, and Safety." *Erkenntnis* 76(1): 121-136.

Smithies, D. 2012. "Mentalism and Epistemic Transparency." *Australasian Journal of Philosophy* 90(4): 723-741.

Sosa, E. 2009. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume II*. Oxford: Oxford University Press.

———. 2015. *Judgment and Agency*. Oxford: Oxford University Press.

Soteriou, M. 2013. *The Mind's Construction: The Ontology of Mind and Mental Action.* Oxford: Oxford University Press.

Srinivasan, A. 2015. "Are We Luminous?" *Philosophy and Phenomenological Research* 90(2): 294-319.

Stalnaker, R. "Luminosity and the KK Thesis." In *Knowledge and Conditionals: Essays on the Structure of Inquiry*, 31-48. Oxford: Oxford University Press, 2019.

Vanrie, W. 2020. "Some Problems with the Anti-Luminosity-Argument." *Pacific Philosophical Quarterly* 101(3): 538-559.

Vogel, J. 2010. "Luminosity and Indiscriminability." *Philosophical Perspectives* 24(1): 547-572.

Weber, Z., & Omori H. 2019. "Observations on the Trivial World." *Erkenntnis* 84: 975-994.

Williamson, T. 1996. "Cognitive Homelessness." *The Journal of Philosophy* 93(11): 554-573.

———. 2000. *Knowledge and its Limits.* New York: Oxford University Press.

———. 2005. "Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge." *The Philosophical Quarterly* 55(219): 213-235.

———. 2008. "Why Epistemology Can't Be Opertationalized." In *Epistemology: New Essays*, edited by Q. Smith, 277-300. Oxford: Oxford University Press.

———. 2021. "The KK principle and rotational symmetry." *Analytic Philosophy* 62(2): 107-124.

Wong, W.-H. 2008. "What Williamson's Anti-Luminosity Argument Really Is." *Pacific Philosophical Quarterly* 89(4): 536-543.