

Please cite the published version from *Pacific Philosophical Quarterly* (<https://doi.org/10.1111/papq.12317>). In that version, I unfortunately forgot to thank Fons Dewulf for making the diagrams.

Some problems with the anti-luminosity-argument*

Wim Vanrie

Abstract. I argue that no successful version of Williamson’s anti-luminosity-argument has yet been presented, even if Srinivasan’s further elaboration and defence is taken into account. There is a version invoking a coarse-grained safety condition, and one invoking a fine-grained safety condition. A crucial step in the former version implicitly relies on the false premise that sufficient similarity is transitive. I show that some natural attempts to resolve this issue fail. Similar problems arise for the fine-grained version. Moreover, I argue that Srinivasan’s defence of the more contentious fine-grained safety condition is also unsuccessful, again for similar reasons.

Timothy Williamson’s anti-luminosity-argument (henceforth: ALA) from *Knowledge and Its Limits* (henceforth: KAIL) aims to show that there are no non-trivial conditions such that, if they obtain, we are always in a position to know that they obtain. The ALA stands to overturn, *inter alia*, several deep-seated philosophical commitments with regards to the mental and its status in epistemology and the philosophy of mind¹. Since it is crucial to know whether it succeeds, the ALA has generated much critical discussion. In her paper *Are We Luminous?* (henceforth: AWL), Amia Srinivasan provides an important further elaboration of the argument, and defends it against past criticism. Since then, there seems to have been a tacit consensus that the combination of Williamson’s original presentation and Srinivasan’s further elaboration yields a more or less ironclad version of the ALA².

In this paper, I argue that this conclusion is too quick. Srinivasan distinguishes two versions of the ALA – both of which are present in Williamson’s original account – one invoking a coarse-grained safety condition on knowledge, and one invoking a fine-grained safety condition. Neither version is successful. The problem with the coarse-grained version is that a crucial step in the argument seems to implicitly rely on the false premise that sufficient similarity is transitive. I present some natural attempts at avoiding this reliance and show that they fail, leaving us with a gap in the argument that remains to be filled. Next, I show that the fine-grained version faces similar problems. Finally, I argue that Srinivasan’s defence of the more contentious fine-grained safety condition is also unsuccessful, for similar reasons.

I do not claim to show that there is no way to repair the ALA. My argument advances our understanding of the (quite intricate) structure of both versions of the ALA, revealing hitherto unnoticed gaps that need to be filled. The burden is on proponents of the ALA to provide us with a successful and compelling version of the argument.

Let us start by rehearsing the core structure of the argument³. A condition R is *luminous* if and only if ‘For every case α , if in α [R] obtains, then in α one is in a position to know that [R] obtains’⁴ (KAIL, 95). The ALA departs from the following scenario (AWL, 296)⁵:

Cold Morning. S wakes up at dawn feeling freezing, very slowly warms up, and feels hot by noon. Throughout the morning S is concentrating sufficiently hard on the question whether she feels cold, such that if she is in a position to know that she feels cold then she does indeed know. S's powers of discrimination are limited, and the change from S's feeling cold to hot is so gradual that S 'is not aware of any change in [her cold feelings] over one millisecond' (KAIL, 97). S's confidence that she feels cold gradually diminishes, such that by noon she firmly believes that she no longer feels cold.

We have cases $\alpha_1, \alpha_2, \dots, \alpha_n$ at times t_1, t_2, \dots, t_n , all a millisecond apart. The ALA starts by supposing, for a *reductio*, that luminosity obtains: if S feels cold in α_i , she is in a position to know that she feels cold in α_i . The following principle is presupposed in the description of 'Cold Morning' (KAIL, 95):

(POS-KNOW-1) If in a case α , S is in a position to know that condition R obtains, and S has done what she is in a position to do to decide whether R obtains, then S knows that R obtains.

Thus, we get:

(LUM) If S feels cold in α_i , then S knows that she feels cold in α_i .

The second premise for the *reductio* is a margin-for-error principle:

(MAR) If S knows that she feels cold in α_i , then S feels cold in α_{i+1} .

(LUM) and (MAR) imply:

(SOR) If S feels cold in α_i , then she feels cold in α_{i+1} .

Thus, (LUM) and (MAR) cannot both obtain. For the ALA to be successful, it needs to be shown that we should reject (LUM), not (MAR). (MAR) is supported by invoking a safety condition on knowledge. There are two versions of the ALA: one invoking a coarse-grained safety condition, and one invoking a fine-grained safety condition. I discuss the coarse-grained version first.

The coarse-grained version of the ALA

The first version of the ALA invokes the following coarse-grained safety condition:

(BELIEF-SAFETY) If in a case α , S knows that condition R obtains, then in any sufficiently similar case β in which S believes that R obtains, it is true that R obtains⁶.

Williamson points out that (BELIEF-SAFETY) by itself does not establish (MAR). Even though S knows that she feels cold at α_i , S might not *believe* that she feels cold at α_{i+1} – since her confidence that she feels cold is gradually decreasing – so that we cannot conclude from

(BELIEF-SAFETY) that she feels cold at α_{i+1} (KAIL, 127). As Srinivasan puts it: ‘To derive (MAR) from (BELIEF-SAFETY), what is needed is a principle that connects S’s belief about C^7 in α_i to S’s belief about C in α_{i+1} ’ (AWL, 301)⁸.

A bad proposal is the following principle, considered by (Berker, 2008, 7) (Vogel, 2010, 561) (Cohen, 2010, 720ff.):

(BEL) If S believes that she feels cold in α_i , then S believes that she feels cold in α_{i+1} .

(BEL) is untenable because it is soritical. It implies that S believes she feels cold at noon, which she does not.

Srinivasan and Williamson both advance a more subtle principle. I shall mainly focus on Srinivasan’s presentation⁹:

(BEL*) If in case α_i , S believes that she feels cold, then there exists a sufficiently similar possible case β_{i+1} in which S’s cold-feelings are a phenomenal duplicate of her cold-feelings¹⁰ in α_{i+1} and in which S believes she feels cold.

Given (BEL*), (BELIEF-SAFETY) yields that S feels cold in β_{i+1} . Because β_{i+1} is a phenomenal duplicate of α_{i+1} , S also feels cold in α_{i+1} , establishing (MAR). (BEL*) exploits the fact that ‘to 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’ (AWL, 302). I should note that Berker and Vogel both reject (BEL*) as a way to establish (MAR) (Berker, 2008, 7) (Vogel, 2010, 562), for reasons that have been adequately addressed by Srinivasan (AWL, 299ff.).

Still, as it stands, (BEL*) is *ad hoc*, and must be supported by argument. Why, for instance, could S’s belief that she feels cold in α_i not be an isolated phenomenon, such that S does not believe that she feels cold in *any* sufficiently similar case? Williamson appeals to considerations of indiscriminability without, however, providing a clear account of how these are meant to support his version of (BEL*) (KAIL, 127). Srinivasan is more elaborate: she takes (BEL*) to be grounded in an empirical principle concerning our doxastic behaviour¹¹:

(DOXDIS) If in a case α , S believes that condition R obtains, then for any case β sufficiently similar to α , S has some disposition¹² in β to believe that R obtains.

Srinivasan notes that she means ‘(DOXDIS) only to be able to handle the kind of central cases of belief-dispositions in play here’ (AWL, 304, footnote 20). For instance, S might have come to believe that R obtains by some strange neurological coincidence. In such *recherché* cases, (DOXDIS) may well be false. But since one counterexample to luminosity is sufficient for the ALA, we can focus on well-behaved cases such as ‘Cold Morning’. In the end, then, (MAR) is supported by ‘a plausible safety condition on knowledge together with a plausible empirical hypothesis about the kind of creatures we are—creatures, namely, whose beliefs are structured by certain kinds of dispositions’ (AWL, 295)¹³.

This is the point at which Srinivasan's elaboration of the coarse-grained version of the ALA ends. Still, (BEL*) does not directly follow from (DOXDIS). Since Srinivasan remains silent on this point, we have to try to fill in the dots ourselves. This is given:

- (1) In α_i , S believes that she feels cold,
- (2) α_i and α_{i+1} are sufficiently similar.

To have (BEL*), we must find a case β_{i+1} satisfying:

- (3) In β_{i+1} , S believes that she feels cold,
- (4) β_{i+1} is a phenomenal duplicate of α_{i+1} ,
- (5) β_{i+1} is sufficiently similar to α_i .

Given (1) and (2), we can apply (DOXDIS) to α_i and α_{i+1} to yield:

- (6) In α_{i+1} , S has some disposition to believe that she feels cold.

This disposition in α_{i+1} must yield the requisite β_{i+1} . The underlying thought, upon which Srinivasan does not elaborate, is that dispositions suitably entail nearby belief (AWL, 303, footnote 19). We can encode this in the following principle:

(DISP-BEL) If in a case α , S has some disposition to believe that condition R obtains, then there is a sufficiently similar phenomenal duplicate¹⁴ β of α in which S believes that R obtains.

(DISP-BEL) need not be uncontroversial: why would the dispositions yielded by (DOXDIS) be strong enough to satisfy (DISP-BEL)? Srinivasan claims that dispositions that would not satisfy (DISP-BEL) arise only in 'unusual cases' (AWL, 303, footnote 19). It is not clear to me, however, on which grounds the luminist would be forced to accept this. Below, we will encounter a weaker alternative that is a more reasonable assumption for the anti-luminist to make¹⁵.


Still, even if (DISP-BEL) is accepted, the argument runs into trouble. Given (6), (DISP-BEL) yields a case β_{i+1} which satisfies:

- (3) In β_{i+1} , S believes that she feels cold,
- (4) β_{i+1} is a phenomenal duplicate of α_{i+1} .

What remains to be established, however, is:

- (5) β_{i+1} is sufficiently similar to α_i .

We do have, with the latter thanks to (DISP-BEL):



- (2) α_i and α_{i+1} are sufficiently similar.
- (7) α_{i+1} and β_{i+1} are sufficiently similar.

(2) and (7) do not yield (5): sufficient similarity is not transitive¹⁶. In taking the step from (DOXDIS) to (BEL*) to be straightforward, it seems that Srinivasan has unknowingly invoked the false premise that sufficient similarity is transitive. It is not clear how (5) could be established. One idea is to exploit the fact that β_{i+1} and α_{i+1} are phenomenal duplicates – i.e. (4) – and use:

(DUPL-SIM-1) If case α is sufficiently similar to β , and β is sufficiently similar to γ , where γ is a phenomenal duplicate of β , then α and γ are sufficiently similar.

Given (2), (4) and (7), (DUPL-SIM-1) yields (5), so that we have (BEL*).

(DUPL-SIM-1), however, when combined with (DOXDIS) and (DISP-BEL), yields a soritical principle. Suppose that, in α_i , S is disposed to believe that she feels cold. Apply (DISP-BEL) to yield a case β_i in which S believes that she feels cold, which is a phenomenal duplicate of α_i , and which is sufficiently similar to α_i . (DUPL-SIM-1) yields that α_{i+1} and β_i are sufficiently similar. We can then apply (DOXDIS) to β_i and α_{i+1} to establish that, in α_{i+1} , S is disposed to believe that she feels cold. Thus, we have:

(DISP-SOR) If in case α_i , S has some disposition to believe that she feels cold, then S has some disposition to believe that she feels cold in α_{i+1} .

This implies that S still has some disposition to believe that she feels cold in the final case α_n . We can apply (DISP-BEL) one final time to get the result that there is a sufficiently similar phenomenal duplicate of α_n wherein S believes that she feels cold, which is absurd.

I conclude that (DOXDIS) cannot be used to establish (BEL*). One last consideration helps, I think, to further enhance the problem. Nothing would seem to preclude the possibility of a series of pairwise sufficiently similar cases $\alpha_1, \alpha_2, \dots, \alpha_n$ satisfying the following properties¹⁷:

- (A) All phenomenal duplicates of α_{i+1} are less similar to α_i than α_{i+1} itself.
- (B) Any case β that is less similar to α_i than α_{i+1} , is not sufficiently similar to α_i .

This yields:

- (C) There is only one phenomenal duplicate of α_{i+1} that is sufficiently similar to α_i , namely α_{i+1} itself.

For such a series of cases, the case β_{i+1} yielded by (BEL*) is always α_{i+1} itself, so that (BEL*) implies that such a series satisfies the soritical principle (BEL), which is absurd.

Srinivasan seems to think that (BEL*) is nothing more than an application to ‘Cold Morning’ of the empirical assumption encoded in (DOXDIS)¹⁸. We now know that this is not true: (BEL*) is not necessary for (DOXDIS). The luminist is free to reject (BEL*) while accepting (BELIEF-SAFETY), (DOXDIS), (DISP-BEL), (POS-KNOW-1), and (LUM).

This is not the end of the ALA, however. In fact, (BEL*) is stronger than what is needed. We do indeed need:

- (3) In β_{i+1} , S believes that she feels cold,
- (4) β_{i+1} is a phenomenal duplicate of α_{i+1} ,


What is not needed, however, is (5) – i.e. sufficient similarity of α_i and β_{i+1} . Given (4), (MAR) also follows from:

- (8) If S knows that she feels cold in α_i , then she feels cold in β_{i+1} .

The aim is to establish (8) through an application of (BELIEF-SAFETY), but this need not be an application of (BELIEF-SAFETY) specifically to α_i and β_{i+1} . It is sufficient to have:

- (9) There exists a phenomenal duplicate β_i of α_i that is sufficiently similar to β_{i+1} .
- (10) If S knows that she feels cold in α_i , then S knows that she feels cold in β_i .

If S knows that she feels cold in α_i , (3), (9) and (10) allow us to apply (BELIEF-SAFETY) to β_i and β_{i+1} to yield that S feels cold in β_{i+1} . This establishes (8). Note, also, that this argument no longer requires that α_{i+1} and β_{i+1} are sufficiently similar, so that we can replace (DISP-BEL) with (DISP-BEL*):



(DISP-BEL*) If in a case α , S has some disposition to believe that condition R obtains, then there is a phenomenal duplicate β of α with regards to the underlying parameters relevant for R in which S believes that R obtains.

Whereas (DISP-BEL) constituted a strong claim about the dispositions yielded by (DOXDIS), it seems reasonable to assume that they satisfy (DISP-BEL*).

(9) is entailed by¹⁹:

(DUPL-SIM-2) If case α is sufficiently similar to β , and γ is a phenomenal duplicate of β , then there exists a case δ that is a phenomenal duplicate of α and that is sufficiently similar to γ .

Although it is hard to definitively establish (DUPL-SIM-2), it also seems hard to deny it. β_{i+1} and α_{i+1} are exactly alike as far as the parameters relevant to S's cold-feelings go, and differ in some other parameters. One can then specify a phenomenal duplicate β_i of α_i that differs from α_i in exactly the same way in which β_{i+1} differs from α_{i+1} . This phenomenal duplicate can be expected to be sufficiently similar to β_{i+1} . Put differently: it seems reasonable to assume that sufficient similarity is such that there is a suitable isomorphism between the spheres of sufficient similarity surrounding phenomenal duplicates.

That leaves (10). Here, there is trouble again. If S knows that she feels cold in α_i , (9) yields that she feels cold in β_i . By luminosity, S is in a position to know that she feels cold in β_i . Since β_i is not part of the original series of cases $\alpha_1, \alpha_2, \dots, \alpha_n$ described in 'Cold Morning', however, we can no longer suppose that S has, in β_i , done what she is in a position to do to decide whether she feels cold, so that we cannot apply (POS-KNOW-1)²⁰. Williamson, in the course of his argument, proposes another principle governing the notion of 'being in a position to know', which we might try instead:

(POS-KNOW-2) If in a case α , S is in a position to know that a condition R obtains, then there is a phenomenal duplicate of α in which S knows that condition R obtains.

Thus, we get that there is a phenomenal duplicate of β_i in which S knows that she feels cold. This, however, puts us in the same situation as before: a further principle is required to establish that this phenomenal duplicate is sufficiently similar to β_{i+1} , so that (BELIEF-SAFETY) can be applied, and no such principle seems to be available. Even if (POS-KNOW-2) would yield a *sufficiently similar* phenomenal duplicate of β_i in which S knows that she feels cold, we would still be stuck with the same kind of triangle as above.

Before proceeding to the fine-grained version of the ALA, I must consider one final attempt to mend the ALA, suggested to me by an anonymous referee. The idea is to exploit the possibility of introducing a more stringent notion of similarity – call it *extreme similarity* – such that the following pseudo-transitivity principle holds:

(EXTREME-TRANS) If case α is extremely similar to β , and β is extremely similar to γ , then α and γ are sufficiently similar.

Assume that the cases have been chosen such that we have the stronger:

(2*) α_i and α_{i+1} are extremely similar.

And suppose that the following also holds:

(DISP-BEL-EXTREME) If in a case α , S has some disposition to believe that condition R obtains, then there is an *extremely* similar phenomenal duplicate β of α in which S believes that R obtains.

Then (7) would be replaced by:

(7*) α_{i+1} and β_{i+1} are extremely similar.

(2*), (7*), and (EXTREME-TRANS) then yield that α_i and β_{i+1} are sufficiently similar, so that (BEL*) holds.

The problem with this proposal is that (DISP-BEL-EXTREME) is *ad hoc*. We have seen that (DISP-BEL) is already a concession to the anti-luminist, which is why we replaced it with (DISP-BEL*). (DISP-BEL-EXTREME) is an even more extreme principle than (DISP-BEL), stating not just that the dispositions yielded by (DOXDIS) yield nearby belief; but that they yield *extremely* nearby belief. There is no pressure on the luminist to accept such a principle. We can conclude that there is, so far, no successful coarse-grained version of the ALA.

The fine-grained version of the ALA

The second version of the ALA is regarded by Williamson as the actual ALA (KAIL, 127, 130). The motivation behind it is that luminists usually suppose that luminous conditions are such that their obtaining constitutively depends on one's belief whether they obtain, a claim that Srinivasan encodes as follows (AWL, 308)²¹:

(CON) If S has done what she is in a position to do to decide whether she feels cold, then she believes that she feels cold if and only if she feels cold.

Srinivasan claims that (CON) implies the falsity of (BEL*) (AWL, 308-309). Her argument, however, involves the claim that phenomenal duplicates are *ipso facto* doxastic duplicates (AWL, 309), which seems to rely on a more crude version of (CON), stating that S believes that she feels cold if and only if she feels cold *simpliciter*. In any case, we have the following: take α_m to be the last case in which S believes that she feels cold, so that α_{m+1} is the first case in which S does not believe that she feels cold. Then (CON) entails that S feels cold in α_m and no longer feels cold in α_{m+1} . Now, the coarse-grained version of the ALA essentially depends on there being a phenomenal duplicate β_{m+1} of α_{m+1} in which S still believes that she feels cold, so that (BELIEF-SAFETY) can be applied. Although this is not in direct contradiction with (CON) – since S need not have done what she is in a position to do to decide whether she feels cold in β_{m+1} – it is still in tension with the obtaining of a constitutive connection between one's feeling cold and one's believing whether one feels cold. Such constitutive connections, after all, are precisely invoked to, if not completely eliminate, at least make false belief that one feels cold into a marginal, defective phenomenon. The stronger the supposed constitutive connection is, the harder it will be to establish a coarse-grained version of the ALA, up to a point where such a connection could foreclose a successful coarse-grained version of the ALA altogether²².

In any case, there is a fine-grained version of the ALA, which is meant to do away with the reliance on false belief, so let us see if it fares any better than the coarse-grained version. To set it up, Williamson invokes his somewhat idiosyncratic notion of 'degrees of confidence', which works as follows. One either believes something outright or not, but outright belief and its absence are compatible with a range of degrees of confidence. Whether one believes outright

or not depends on whether one's degree of confidence is above a certain threshold. Importantly, degrees of confidence are not subjective probabilities (KAIL, 99): A sufficiently high degree of confidence in p entails that one believes outright that p , which is not the same as assigning a high probability to p .

To exploit these fine-grained degrees of confidence, (BELIEF-SAFETY) is replaced by²³:

(CONFIDENCE-SAFETY) If in a case α , S knows that condition R obtains, then in any sufficiently similar case β in which S has an at-most-slightly-lower degree of confidence that she is in condition R, it is true that R obtains.

(CONFIDENCE-SAFETY) is stronger than (BELIEF-SAFETY). The crucial difference is this: even if, in a case α , there is *no* nearby case in which one falsely believes that R obtains, (CONFIDENCE-SAFETY) can still preclude knowledge if one's degree of confidence in α is only slightly above the threshold for outright belief and one's degree of confidence in a nearby case is slightly below the threshold for outright belief. As Srinivasan puts it, even 'nearby untrue *almost-belief*' (AWL, 314) suffices to undermine knowledge.

Many philosophers have claimed that (CONFIDENCE-SAFETY) is an unreasonably strong requirement on knowledge²⁴. Srinivasan has replied to these critiques (AWL, 311ff.). I discuss her defence below. Even if (CONFIDENCE-SAFETY) is accepted, however, there remain problems with the fine-grained version of the ALA analogous to those with the coarse-grained version.

In their description of the case, Srinivasan and Williamson assume that 'one's confidence that one feels cold gradually decreases' (KAIL, 97) in such a way that the following holds:

(CONF) If in α_i , S has degree of confidence c that she feels cold, then in α_{i+1} , S has an at-most-slightly-lower degree of confidence c^{-25} that she feels cold.

(CONF) stands to (CONFIDENCE-SAFETY) as the soritical (BEL) stands to (BELIEF-SAFETY): it establishes that S's doxastic attitude throughout the series of cases behaves in such a way that the corresponding safety principle can be applied to consecutive cases. Contrary to (BEL), however, (CONF) is not soritical, precisely because (CONFIDENCE-SAFETY) is fine-grained and thereby compatible with slight variations in S's degree of confidence²⁶.

Given (CONF), (CONFIDENCE-SAFETY) can be straightforwardly applied to any two consecutive cases, immediately yielding (MAR). The question, however, is whether (CONF) is as innocuous a premise as it might appear to be. Recall that it is stipulated that 'throughout the process one thoroughly considers how cold or hot one feels' (KAIL, 97). If feeling cold *is* luminous, and because S is thoroughly considering how cold she feels throughout, we can expect the transition between S's feeling cold and her not feeling cold to be reflected in a concomitant transition in her degree of confidence that she feels cold, in a way that falsifies (CONF). More specifically, if α_m is the last case in which she feels cold, and α_{m+1} the first case in which she does not, (CONF) need not obtain for those cases. In α_{m+1} , S knows that she

has stopped feeling cold, so her degree of confidence that she feels cold will be adjusted accordingly²⁷.

It may be thought that this straightforwardly contradicts what is, for Williamson and many others, the undeniable fact that cases such as α_m and α_{m+1} are indiscriminable. But this all depends on how one characterizes ‘indiscriminability’. As we will see, Srinivasan’s way of fleshing it out in terms of doxastic dispositions is perfectly compatible with the denial of (CONF). One confusion that seems to be at work sometimes is that between knowing, in a case α , that one feels cold, and knowing that one feels cold in α . The luminist is not committed to the luminosity of ‘being in case α ’, ‘feeling cold in case α ’, and the like. Some critics of luminosity nevertheless seem to have such conditions in mind. Vogel, for instance, emphasizes that it is implausible that one would be able to ‘[compare] the two rosters of properties pertaining to [α_m] and [α_{m+1}] respectively, and [recognize] that [α_m] is different from [α_{m+1}]’ (Vogel, 2010, 564). But luminists are not committed to this. They need not claim that one can know, *of* sufficiently similar (indiscriminable) cases, that they are different. Rather, luminists claim that one can know things that are different in sufficiently similar cases, which is not the same thing.

Srinivasan points out that Berker, who opposes the ALA, himself says that (CONF) ‘seems indisputable, given the description of the situation at hand’ (Berker, 2008, 12) (AWL, 310). She even says that Berker ‘proposes’ (AWL, 310) (CONF) himself, but this is somewhat misleading. Rather than proposing it, he extracts it from Williamson’s description of the case (Berker, 2008, 11). I take Berker’s acceptance of (CONF) to be only for the sake of argument, since he focuses his criticism on (CONFIDENCE-SAFETY). Indeed, his remark seems somewhat tongue-in-cheek: of course (CONF) is indisputable *given the description of the situation at hand*, since it is included in that description. The whole question is whether the luminist has to accept that description. I expect Berker to be open to resisting (CONF) as an alternative way of resisting the ALA. Indeed, some of his remarks towards the end of his paper seem to be moving in this direction (Berker, 2008, 16ff.). (CONF) cannot simply be assumed, it must be established.

In fact, notwithstanding her remarks about Berker, Srinivasan seems to agree, since she proposes a fine-grained version of the ALA that does not rely on (CONF). In the same way that (BEL) is replaced with the more sophisticated (BEL*), she replaces (CONF) with (AWL, 310):

(CONF*) If in α_i , S has degree of confidence c that she feels cold, there exists a sufficiently similar possible case β_{i+1} in which S’s cold-feelings are a phenomenal duplicate of her cold-feelings in α_{i+1} and in which S has an at-most-slightly-lower degree of confidence c^- that she feels cold.

Given (CONF*), (CONFIDENCE-SAFETY) yields (MAR). Again, however, (CONF*) must be supported by further argument. To this end, Srinivasan advances a fine-grained version of (DOXDIS) (AWL, 311):

(DOXDIS-CONF)²⁸ If in a case α , S believes with degree of confidence c that condition R obtains, then in any case β sufficiently similar to α , S has some disposition in β to believe with an at-most-slightly-lower degree of confidence c^- that R obtains.

Again, this is meant to be ‘a plausible empirical premise about our dispositions to believe similarly in similar situations’ (AWL, 311). As before, however, (DOXDIS-CONF) does not directly yield (CONF*), and Srinivasan does not elaborate. What is given, is:

- (1) In α_i , S has degree of confidence c that she feels cold,
- (2) α_i and α_{i+1} are sufficiently similar.

For the fine-grained argument, we must find a case β_{i+1} satisfying:

- (3) In β_{i+1} , S has an at-most-slightly-lower degree of confidence c^- that she feels cold than in α_i ,
- (4) β_{i+1} is a phenomenal duplicate of α_{i+1} ,
- (5) β_{i+1} is sufficiently similar to α_i .

Given (1) and (2), we can apply (DOXDIS-CONF) to α_i and α_{i+1} :

- (6) In α_{i+1} , S is disposed to believe with an at-most-slightly-lower degree of confidence c^- that she feels cold.

To get the requisite β_{i+1} , a fine-grained version of (DISP-BEL) can be used:

(DISP-CONF) If in a case α , S has some disposition to believe with degree of confidence c that condition R obtains, then there is a case β , sufficiently similar to α , which is a phenomenal duplicate of α with regards to the underlying parameters relevant for R, in which S believes with degree of confidence c that R obtains.

As with (DISP-BEL), it is not clear why the luminist should be compelled to accept (DISP-CONF). Below, I introduce a fine-grained version of the more plausible alternative (DISP-BEL*). Again, however, there is trouble even with (DISP-CONF). As before, (DISP-CONF) yields a case β_{i+1} which satisfies (3) and (4), but not (5). What we have is (2) and:

- (7) α_{i+1} and β_{i+1} are sufficiently similar.

We could again try:

(DUPL-SIM-1) If case α is sufficiently similar to β , and β is sufficiently similar to γ , where γ is a phenomenal duplicate of β , then α and γ are sufficiently similar.

In the fine-grained case, (DUPL-SIM-1) has no soritical implications: it only yields that, if S is disposed, in α_i , to believe with degree of confidence c that she feels cold, then S is disposed, in α_{i+1} , to believe with an at-most-slightly-lower degree of confidence c^- that she feels cold, which is unproblematic. There is still a serious problem with (DUPL-SIM-1), however. Let

α_{i+1}^2 be a sufficiently similar phenomenal duplicate of α_{i+1} . Applying (DUPL-SIM-1) to α_i , α_{i+1} , and α_{i+1}^2 yields that α_i is sufficiently similar to α_{i+1}^2 . Now let α_{i+1}^3 be a sufficiently similar phenomenal duplicate of α_{i+1}^2 . Applying (DUPL-SIM-1) to α_i , α_{i+1}^2 , and α_{i+1}^3 yields that α_i is sufficiently similar to α_{i+1}^3 . And so on. The result is that, for any chain $\alpha_{i+1} = \alpha_{i+1}^1, \alpha_{i+1}^2, \dots, \alpha_{i+1}^n$ of pairwise sufficiently similar phenomenal duplicates of α_{i+1} , each member of that chain is sufficiently similar to α_i , which is absurd: in α_{i+1}^n , S might well be a body in a vat with suitable temperature.

As in the coarse-grained version, however, (CONF) is stronger than what is required. It is sufficient to have:

- (8) There exists a phenomenal duplicate β_i of α_i that is sufficiently similar to β_{i+1} .
- (9) S's degree of confidence c^- that she feels cold in β_{i+1} is at-most-slightly-lower than her degree of confidence that she feels cold in β_i .
- (10) If S knows that she feels cold in α_i , then S knows that she feels cold in β_i .

(MAR) can then be established by applying (CONFIDENCE-SAFETY) to β_i and β_{i+1} . Once again, the sufficient similarity of α_{i+1} and β_{i+1} is no longer required, so that we can replace (DISP-CONF) with (DISP-CONF*):

(DISP-CONF*) If in a case α , S has some disposition to believe with degree of confidence c that condition R obtains, then there is a phenomenal duplicate β of α with regards to the underlying parameters relevant for R in which S believes with degree of confidence c that R obtains.

As before, (8) is entailed by (DUPL-SIM-2).

Next up is (9). Here, there is a new problem. In effect, (9) states that the difference between S's degrees of confidence in β_i and β_{i+1} is small enough so that (CONFIDENCE-SAFETY) can be applied to them. S's degree of confidence in β_{i+1} is c^- . There is nothing to prevent, however, that S's degree of confidence that she feels cold in β_i is higher than in α_i , say $d > c$, in such a way that the difference between d and c^- is too large to apply (CONFIDENCE-SAFETY), whereby (9) is false. I can think of only one – somewhat contrived – way of overcoming this. We can apply the same principles to the pair of cases β_i and β_{i+1} , as we did to α_i and α_{i+1} . This yields sufficiently similar cases γ_i and γ_{i+1} such that: (3') in γ_{i+1} , S has an at-most-slightly-lower degree of confidence d^- that she feels cold than in β_i ; (8') γ_i is a phenomenal duplicate of β_i that is sufficiently similar to γ_{i+1} . Again, it remains to be established that: (9') S's degree of confidence d^- that she feels cold in γ_{i+1} is at-most-slightly-lower than her degree of confidence e that she feels cold in γ_i . Notice, however, that if (9') is also false, this must be because $e > d$. Thus, we can repeat the same procedure, which will either yield the right sort of cases, or a further phenomenal duplicate of γ_{i+1} in which S's degree of confidence is strictly greater than e . If we now assume that the spectrum of degrees of confidence available to a human subject is discrete, it follows from the fact that there is a maximal degree of confidence that a finite amount of iterations will yield a pair of cases ζ_i and ζ_{i+1} which satisfy the analogous versions of (3), (8), and (9). Whether this assumption is

problematic, I cannot assess here. In any case, if it is not granted, I see no way of establishing (9), which spells trouble for the ALA.

Even if (9) is granted, (10) faces the same problem as before. Luminosity only yields that, in β_i , S is in a position to know that she feels cold, but we cannot suppose that S has, in β_i , done what she is in a position to do to know whether she feels cold. Again, (POS-KNOW-2) offers no way out. I conclude that we have not yet been presented with a successful fine-grained version of the ALA either, even if (CONFIDENCE-SAFETY) is accepted.

Moreover, there are problems with Srinivasan's defence of (CONFIDENCE-SAFETY).

Srinivasan's defence of (CONFIDENCE-SAFETY)

Williamson does not provide a sustained argument for (CONFIDENCE-SAFETY). As pointed out, it has been the subject of much criticism, to which Srinivasan attempts to reply²⁹. I now argue that her defence of (CONFIDENCE-SAFETY) is unsuccessful³⁰. We can focus on her 'Glass Half Empty' case (AWL, 313):

Glass Half Empty. Henrietta likes watching full glasses slowly drain until they are empty. The confidence threshold for outright belief is 80%, and, in normal conditions and when she is paying close attention, Henrietta believes of a given glass that it is at least half empty if and only if it is indeed at least half empty. Henrietta's confidence that the glass is at least half empty remains near 0% until the glass is almost half empty, then steeply increases to 80% just as the glass is exactly half empty, increasing to and remaining at 100% shortly thereafter. The only proposition that Henrietta entertains as the glass empties is q , *that the glass is at least half empty*.

(CONFIDENCE-SAFETY) implies that Henrietta does not know that q when the glass is exactly half empty. Srinivasan endorses this conclusion on the basis of 'a common picture of the relationship between confidence and practical reasoning' (AWL, 313):

'According to that picture, one's confidence in a proposition p is a measure of one's willingness to rely on p as a premise in practical reasoning; one believes outright when one's willingness to rely on p crosses a certain threshold. If so, then a confidence just short of outright belief in p will yield some cases in which one uses p in one's practical reasoning despite not believing p outright. When the glass is only, say, 49% empty, Henrietta's 79% confidence that the glass is at least half empty means that she has some tendency to use that false proposition as a premise in her practical decision-making. If she were making several q -relevant but independent decisions at the same time, we could expect to see Henrietta acting on q when it is false. This nearby willingness to act on q when it is untrue intuitively undermines Henrietta's claim to know q in the case where q has just become true' (AWL, 313-314).

As is clear from the last sentence, Srinivasan aims to bolster (CONFIDENCE-SAFETY) by invoking yet another safety condition, which we can formulate as follows:

(PRACTICAL-SAFETY) If in a case α , S knows that condition R obtains, then in any sufficiently similar case β in which S is willing to act on R's obtaining, it is true that R obtains.

Henrietta does not know that q when the glass is exactly half empty, because there are nearby cases in which she is willing to act on the false premise q .

As will be expected by now, (PRACTICAL-SAFETY) does not by itself undermine Henrietta's claim to know that q when the glass is exactly half empty, since there must *be* a sufficiently similar case wherein she is willing to act on the false premise q . The crucial question is: when does one count as so willing to act? Call α_i the case in which the glass is exactly 50% empty. In α_i , Henrietta believes with 80% confidence that q . Just before α_i , we have the sufficiently similar case α_{i-1} , in which the glass is only 49% empty and in which Henrietta is 79% confident that it is at least half empty. Srinivasan wishes to apply (PRACTICAL-SAFETY) to α_i and α_{i-1} , so she must establish that Henrietta indeed counts as willing to act on q in α_{i-1} . The underlying idea is that, if one's degree of confidence in q is just below the threshold for outright belief, then one counts as 'willing to act on q '. To support this, Srinivasan writes that 'Henrietta's 79% confidence that the glass is at least half empty means that she has some tendency to use that false proposition as a premise in her practical decision-making. If she were making several q -relevant but independent decisions at the same time, we could expect to see Henrietta acting on q when it is false' (AWL, 314). But this is too weak. On this line of thought, 70% confidence also entails that Henrietta has 'some tendency to use q as a premise in her practical decision-making', since exactly the same sort of consideration applies: if Henrietta had 70% confidence, or even 40% confidence, and she had to make several q -relevant but independent decisions at the same time, we could expect to see her act on q when it is false. So Henrietta counts as 'willing to act on q ' even in those cases. But 70% confidence does not undermine knowledge. Thus, Srinivasan's notion of 'some tendency to use a proposition in one's practical decision-making' is too liberal: one's willingness to act on q in a nearby case in *this* sense is not sufficient to undermine knowledge.

What is required, is a more robust notion of what it means to be 'willing to act' on a premise in such a way that it does undermine knowledge. This seems to be a natural proposal:

(WILL-ACT) In a case α , S is willing to act on the proposition that condition R obtains if there is a sufficiently similar phenomenal duplicate³¹ β in which S acts on the proposition that condition R obtains.

In effect, the idea is to conceive of willingness to act on a premise as a disposition to act on that premise which is actualized in other cases. The question then becomes: how strong must this disposition be in order to undermine knowledge, as stated in (PRACTICAL-SAFETY)? My proposal is that it must be strong enough so as to be actualized, not just in any remote case, but in a *sufficiently similar* case³².

Given (WILL-ACT), we have to establish:

(1) There exists a case β_{i-1} that is sufficiently similar to α_{i-1} and in which Henrietta acts on q .

What has to yield (1), is the fact that Henrietta is, in α_{i-1} , 79% confident that q , i.e. that her degree of confidence is just below the threshold for outright belief. A natural attempt is:

(ACT-CONF) If in a case α , S has a degree of confidence that is at-most-slightly-lower than the threshold for outright belief that R obtains, then there is a sufficiently similar phenomenal duplicate β in which S acts on R's obtaining.

(ACT-CONF) yields (1), so that (WILL-ACT) yields that Henrietta counts as willing to act on q in α_{i-1} , as required. Srinivasan, however, provides no compelling reason to accept (ACT-CONF). She writes: 'A confidence just short of outright belief in p will yield some cases in which one uses p in one's practical reasoning despite not believing p outright' (AWL, 314). But 'some cases' is, again, too weak. What is needed, is a sufficiently similar case. It seems reasonable to think that a degree of confidence *above* the threshold for outright belief yields such a case, but I can see no good reason to think that a degree of confidence *below* that threshold also yields such a case. Indeed, it seems natural to think that the threshold for one's counting as willing to act coincides with the threshold for outright belief. Srinivasan writes:

'At the point at which the glass is just less than half empty, Henrietta is 79% confident that the glass is at least half empty; very confident indeed. Does this not undermine Henrietta's claim to know q just a moment later, when the glass is exactly half empty?' (AWL, 313).

I do not think that a degree of confidence that falls below the threshold for outright belief is adequately described as 'very confident indeed'. Imagine the following conversation: 'Do you believe that your partner is cheating on you?' – 'No, but I am very confident that she is'. To me, this sounds very infelicitous indeed. In any case, no compelling reason has been given to accept that a 79% degree of confidence that falls below the threshold of outright belief yields the sort of willingness to act that is required for (PRACTICAL-SAFETY). I conclude, then, that Srinivasan's defence of (CONFIDENCE-SAFETY) fails, so that the fine-grained version of the ALA is in even more trouble.

Of course, I have discussed just one natural attempt of fleshing out the relevant notion of 'willingness to act'. Given that it fails, the burden, lies with the proponents of (CONFIDENCE-SAFETY) to provide a better defence of it.

Department of Philosophy and Moral Sciences
Ghent University

NOTES

* My thanks go out to Amia Srinivasan for helpful comments on an earlier version of this paper, as well as to several anonymous referees for their fruitful remarks. I also thank Maarten Van Dyck for several discussions.

¹ See (Hawthorne & Srinivasan, 2013) (Srinivasan, 2015b) for some discussion of the upshot of the ALA.

² The survey in (McGlynn, 2014, chapter 7) – which incorporates Srinivasan’s paper – is a good illustration of this. I know of no direct criticism of Srinivasan’s account. Those who still wish to defend luminosity do so in ways that are orthogonal to it. One prominent such strategy is to invoke the (Chalmers, 2003) notion of ‘direct phenomenal concepts’ (Barz, 2017) (Duncan, 2018). (Jenkins, 2018) accepts the ALA, but argues that luminosity still obtains for certain mental processes and events (rather than mental conditions).

³ My presentation is indebted to Berker’s helpful regimentation of the argument (Berker, 2008), which was also adopted by Srinivasan.

⁴ ‘A case depends on a subject (referred to by ‘one’), a time (referred to by the present tense), and a possible world. [...] The domain of cases will be taken to include counterfactual as well as actual possibilities. [...] Conditions are coarsely individuated by the cases in which they obtain: they are identical if they obtain in exactly the same cases’ (KAIL, 94). We need not worry about the precise nature of conditions and cases.

⁵ I use Srinivasan’s formulation. Whereas Williamson designates the subject of a case by ‘one’, Srinivasan uses ‘S’. Nothing important hangs on this, and I will freely switch between the two.

⁶ Williamson: ‘If one believes p truly in a case α , one must avoid false belief in other cases sufficiently similar to α in order to count as reliable enough to know p in α ’ (KAIL, 100). Compare also (9) at (KAIL, 128).

⁷ C is the condition of one’s feeling cold.

⁸ See also (Blackson, 2007, 402f.) (Wong, 2008, 539) (Berker, 2008, 6ff.) (Ramachandran, 2009, 663) (Cohen, 2010, 728) (Greenough, 2012, 333) (Ramachandran, 2012, 132) (Smithies, 2012, 725-726).

⁹ Williamson’s version of (BEL*) is (8) at (KAIL, 127). It is a straightforward exercise to adapt my discussion to Williamson’s account.

¹⁰ If one’s cold feelings in one case are a phenomenal duplicate of one’s cold feelings in another, one feels cold in the one if and only if one feels cold in the other. Compare (7) from (KAIL, 127). I call two cases *phenomenal duplicates* if S’s cold-feelings in the one are a phenomenal duplicate of her cold-feelings in the other.

¹¹ I have altered Srinivasan’s formulation, which does not accord with Williamson’s case/condition terminology.

¹² I adopt this somewhat unclear notion of ‘having some disposition’ from Srinivasan. Below, we will encounter reasons why it would have to be fleshed out.

¹³ Srinivasan is right to say that ‘what is at stake in the debate about anti-luminosity is a certain vision of what kind of creatures we are, empirically speaking’ (AWL, 295), as comes out in Williamson’s appeal to our limited powers of discrimination (KAIL, 13, 127). This raises suspicion about objections to the ALA based on stipulated cognitive make-ups – e.g. (Leitgeb, 2002, 216) (Weatherson, 2004, 379) – as also noted by (Ramachandran, 2009, 668). The question is: are there conditions such that *we*, with the cognitive apparatus we actually possess, are always in a position to know that they obtain when they obtain?

¹⁴ Here, I am expanding the notion of a phenomenal duplicate to arbitrary conditions R .

¹⁵ A further objection has been suggested to me by an anonymous referee. The idea is that S may have some disposition to believe that she feels cold, but that this disposition may be such that a change in her phenomenal state is required for it to be actualized, on the following model: ‘I have some disposition to walk to the fridge and get a snack; but that disposition might only be triggered if my feelings of hunger become somewhat stronger than they currently are’ (quoted from the referee’s report). In that case, there would be *no* phenomenal duplicate in which S believes that she feel cold. This would spell further trouble for (DISP-BEL) and for the ALA. I leave it as a task for the anti-luminist to address this concern, which presumably requires going into more detail about the relevant notion of ‘disposition’.

¹⁶ The problem was noted by (Vogel, 2010, 569, endnote 51). Related observations are made by (Zardini, 2012, 398f.).

¹⁷ Intuitively, the idea is that the successive cases are chosen such that they are as far apart as they can be while still being pairwise sufficiently similar.

¹⁸ Note, in this regard, the ‘thus’ in her summary of the coarse-grained argument: ‘Suppose that S is in a condition C , but in what we might call a ‘liminal’ case of it. That is, there is an extremely similar case to the one she is in which is not a case of C . Imagine that S believes she is in C ; is this belief knowledge? It seems not. For in the very similar non- C case, S has the disposition to believe she is in C . Thus, her true belief that she is in C is rendered

unreliable by a nearby untrue belief that she is in C' [my emphasis] (AWL, 305-306). Srinivasan seems to have underestimated the difficulty of unpacking this 'thus'.

¹⁹ We can see (DUPL-SIM-2) as 'completing the square' as in the figure above by providing the requisite case β_i .

²⁰ An anonymous referee pointed out that the following safety condition would also do the trick: (POSITION-SAFETY) If in a case α , S is in a position to know that condition R obtains, then in any sufficiently similar case β in which S believes that R obtains, it is true that R obtains. We can then simply apply (POSITION-SAFETY) to β_i and β_{i+1} to yield (10). The problem with this is that (POSITION-SAFETY) entails, by itself, the falsity of luminosity. Consider all cases in which S is just barely below the threshold for feeling cold. Since we are considering *all* such cases, there is one in which S nevertheless believes that she feels cold, call it β . Because S just barely does not feel cold in β , there is a sufficiently similar case α in which S does feel cold. Luminosity would entail that S is in a position to know that she feels cold in α . But (POSITION-SAFETY) implies that S cannot be in a position to know that she feels cold in α , since β is a sufficiently similar case in which she falsely believes that she feels cold. Consequently, the luminist will automatically reject (POSITION-SAFETY). Of course, it may be that there are strong independent grounds for (POSITION-SAFETY). If so, this would constitute a (surprising and interesting) alternative version of the ALA. Much will hang on how exactly one construes the state of 'being in a position to know'. Exploring this further, however, is work for another paper.

²¹ For versions of such a constitutivity thesis, see (Heal, 2001, 4) (Weatherston, 2004, 379) (Berker, 2008, 9) (Coliva, 2009, 367). Some go even further and claim an equivalence between a condition's obtaining and one's *knowing* that it obtains (Leitgeb, 2000, 204) (Hossack, 2002, 174). In the present context, this is question-begging.

²² The crude version of (CON) mentioned above would be an example.

²³ This is Srinivasan's formulation, slightly altered. See also (Berker, 2008, 11). Williamson: 'If one believes outright to some degree that a condition C obtains, when in fact it does, and at a very slightly later time one believes outright on a very similar basis to a very slightly lower degree that C obtains, when in fact it does not, then one's earlier belief is not reliable enough to constitute knowledge' (KAIL, 101).

²⁴ E.g. (Leitgeb, 2002, 202ff.) (Blackson, 2007, 204-205) (Berker, 2008, §4) (Ramachandran, 2009, §VI) (Cohen, 2010, 726) (Zardini, 2012, §6ff.).

²⁵ Throughout, I will use the affix '-' to indicate 'at-most-slightly-lower'.

²⁶ Although (BELIEF-SAFETY) is compatible with *some* such variations, it is not compatible with variations that cross the threshold for outright belief.

²⁷ In this way, Williamson's claim that S 'is not aware of any change in [her cold feelings] over one millisecond' (KAIL, 97), repeated by Srinivasan (AWL, 296), can be read in a way that is question-begging.

²⁸ I have altered Srinivasan's formulation, which does not accord with Williamson's terminology.

²⁹ Since my discussion does not require it, I shall not discuss those criticisms here. For some examples, see (Leitgeb, 2002, 202ff.) (Berker, 2008, §4) (Ramachandran, 2009, 667ff.) (Cohen, 2010, 725ff.) (Zardini 2012, §6). See also (McGlynn, 2014, §7.2) for relevant discussion.

³⁰ Note that, if the luminist rejects (CONFIDENCE-SAFETY), it becomes open to her to accept (CONF), discussed above. The ALA can be independently resisted on both fronts.

³¹ The demand that β is a phenomenal duplicate is necessary to preserve the falsity of the premise on which S acts.

³² Sufficiently similar to the case in which one is willing to act, not to the case for which knowledge is claimed. The claim is not that knowledge is only undermined by a nearby case in which one *does* act on a false premise. Knowledge is still conceived as being undermined by a nearby willingness to act on a false premise, but that nearby willingness is now understood, in turn, as nearby acting on that premise. Sliding over some details, we could say that knowledge is undermined by nearby nearby acting on a false premise.

REFERENCES

- Barz, W. (2017). 'Luminosity Guaranteed,' *Pacific Philosophical Quarterly* 98, pp. 480-496.
Berker, S. (2008). 'Luminosity Regained,' *Philosopher's Imprint* 8, pp. 1-22.
Blackson, T. A. (2007). 'On Williamson's Argument for (I_i) in His Anti-Luminosity Argument,' *Philosophy and Phenomenological Research* 74, pp. 397-405.

- Chalmers, D. (2003). 'The content and epistemology of phenomenal belief,' in Q. Smith and A. Jolic (eds) *Consciousness: New Philosophical Perspectives*. (Oxford: Oxford University Press) pp. 220-272.
- Cohen, S. (2010). 'Luminosity, Reliability, and the Sorites,' *Philosophy and Phenomenological Research* 81, pp. 718-730.
- Coliva, A. (2009). 'Self-knowledge and commitments,' *Synthese* 171, pp. 365-375.
- Duncan, M. (2018). 'What it's like to have a cognitive home,' *European Journal of Philosophy* 26, pp. 66-81.
- Greenough, P. (2012). 'Discrimination and self-knowledge,' in D. Smithies and D. Stoljar (eds) *Introspection and Consciousness*. (Oxford: Oxford University Press) pp. 329-349.
- Hawthorne, J. and Srinivasan, A. (2013). 'Disagreement Without Transparency: Some Bleak Thoughts,' in D. Christensen and J. Lackey (eds) *The Epistemology of Disagreement: New Essays*. (Oxford: Oxford University Press) pp. 9-30.
- Heal, J. (2001). 'On First-Person Authority,' *Proceedings of the Aristotelian Society* 102, pp. 1-19.
- Hossack, K. (2002). 'Self-Knowledge and Consciousness,' *Proceedings of the Aristotelian Society* 102, pp. 163-181.
- Jenkins, D. (2018). 'Luminosity in the stream of consciousness,' *Synthese*. doi:10.1007/s11229-018-1801-0.
- Leitgeb, H. (2002). 'Review of Timothy Williamson, *Knowledge and its Limits*,' *Grazer Philosophische Studien* 65, pp. 195-205.
- McGlynn, A. (2014). *Knowledge First?* (London: Palgrave Macmillan)
- Ramachandran, M. (2009). 'Anti-Luminosity: Four Unsuccessful Strategies,' *Australasian Journal of Philosophy* 84, pp. 659-673.
- Ramachandran, M. (2012). 'The KK-Principle, Margins for Error, and Safety,' *Erkenntnis* 76, pp. 121-136.
- Smithies, D. (2012). 'Mentalism and Epistemic Transparency,' *Australasian Journal of Philosophy* 90, pp. 723-741.
- Srinivasan, A. (2015). 'Are We Luminous?,' *Philosophy and Phenomenological Research* 90, pp. 294-319.
- Srinivasan, A. (2015). 'Normativity Without Cartesian Privilege,' *Philosophical Issues* 25, pp. 273-299.
- Vogel, J. (2010). 'Luminosity and Indiscriminability,' *Philosophical Perspectives* 24, pp. 547-572.
- Weatherson, B. (2004). 'Luminous Margins,' *Australasian Journal of Philosophy* 82, pp. 373-383.
- Williamson, T. (2000). *Knowledge and its Limits*. (Oxford: Oxford University Press)
- Wong, W.-H. (2008). 'What Williamson's Anti-Luminosity Argument Really Is,' *Pacific Philosophical Quarterly* 89, pp. 536-543.
- Zardini, E. (2012). 'Luminosity and Vagueness,' *Dialectica* 66, pp. 375-410.