



DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD

Luminosity Regained

The Harvard community has made this article openly available.
Please share how this access benefits you. Your story matters.

Citation	Berker, Selim. 2008. Luminosity regained. <i>Philosophers' Imprint</i> 8(2): 1-22, http://hdl.handle.net/2027/spo.3521354.0008.002 (accessed February 26 2009).
Published Version	http://quod.lib.umich.edu/cgi/t/text/text-id?c=phimp;rgn=main;idno=3521354.0008.002
Accessed	April 16, 2018 6:52:31 PM EDT
Citable Link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:2643118
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

(Article begins on next page)

Luminosity Regained

Selim Berker

Harvard University

© 2008 Selim Berker

<www.philosophersimprint.org/008002/>

IN THE ONGOING SKIRMISHES between internalists and externalists in epistemology, one notion that seems particularly resistant to an externalist treatment is that of rationality. Perhaps I can fail to know or even be justified in believing that *p* without having epistemic access to whatever it is that deprives me of knowledge or justification. However, can it be *irrational* for me to believe that *p* when part of what makes that the case is outside my cognitive grasp, when from my first-person perspective everything indicates that it would be perfectly appropriate for me to so believe? Rationality seems to involve doing the best one can with what one has, and to posit an externalist standard of rationality according to which it can be irrational for one to believe a given proposition despite one's not being able, even upon reflection, to realize that this is so at best sounds like a strained use of the word 'irrational', and at worst appears to be changing the subject.

Not so, according to Timothy Williamson. In *Knowledge and Its Limits*,¹ Williamson provides a sorites-like argument for the conclusion that *no* plausible standard of rationality meets the constraint that its demands are always accessible to the subject. In Williamsonian parlance, the fact that it would be irrational of one to believe that *p* is not *luminous*, for that fact can obtain without one's being in a position to know that it obtains. So not only is an externalist notion of rationality fully coherent, but it is the only game in town: even if we posit two notions of rationality, one of which (rationality₁) is claimed to be externalist and the other of which (rationality₂) is claimed not to be, the latter will be just as susceptible to Williamson's argument as the former. Therefore those who cling to the idea that rationality is inherently internalistic should cast off their prejudices and learn to love the bomb; sometimes we lack access to the demands of rationality, but such is our lot, and to hope for things to be otherwise is to quixotically yearn after an impossible cognitive standard.

Williamson's anti-luminosity argument has further consequences. If sound, not only would it show that it can be irrational for one to believe a given proposition without one's being in a position to know

1. Williamson (2000). All page references in the text are to this book.

that it is, but it would also establish that one *can be in pain* without one's being in a position to know that one is; that one *can seem to see* a certain color without one's being in a position to know that one does; that two words *can mean the same thing* in one's idiolect without one's being in a position to know that they do; and that a given proposition *can be part of one's evidence* without one's being in a position to know that it is. In other words, the argument would show that one's current mental life, the meanings of one's words, the extent of one's evidence, and the dictates of rationality are all non-luminous — that each can, at least in principle, be epistemically inaccessible to a given subject. If Williamson is right, then we are (as he puts it) “cognitively homeless”: there is no substantive domain of mental or semantic or normative facts to which we have guaranteed access, no subportion of our mental or semantic or normative lives within which everything lies open to view.

However, we need not accept Williamson's argument. Luminosity is a kind of *epistemic* privileged access: it involves a subject's always being in a position to know that a given fact obtains, whenever it does obtain. I will argue that Williamson's argument 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. The exact nature of this connection will depend on the version of Williamson's argument that one is considering: as we shall see, there are two versions of the anti-luminosity argument, depending on whether one defends the argument's crucial premise in terms of “all or nothing” belief or in terms of degrees of belief. But in either case, the thesis at the level of belief that Williamson must deny in order to secure his results at the level of knowledge is one that is independently plausible, and one to which defenders of luminosity will readily help themselves; to simply *assume* its falsity would beg the main question at issue. Only those who follow Williamson in his radical claims about how, even after ideal reflection, our beliefs about our mental lives, the meanings of our words, and the demands of rationality can swing free

from the truth of these matters need accept his equally radical claims about our inevitable cognitive homelessness.

1. Williamson's Core Argument

In Williamson's terminology, a condition *c* is *luminous* if and only if the following holds:

- (*) For every case α , if in α condition *c* obtains, then in α one is in a position to know that *c* obtains.

What Williamson calls *cases* are what Lewis, following Quine, calls centered possible worlds: possible worlds with a designated subject and a designated time. Williamsonian *conditions*, on the other hand, are in effect centered states-of-affairs: a condition either obtains or fails to obtain in a given case, and each condition can be specified by a that-clause in which the pronoun “one” refers to the case's designated subject and the present tense refers to the case's designated time. So, for example, the condition *that one has hands* obtains in a case α if and only if in α the subject of α has hands at the time of α . The expression “in a position to know” in (*) is potentially obscure, but for our purposes we only need the following: according to Williamson, “if one is in a position to know *p*, and one has done what one is in a position to do to decide whether *p* is true, then one does know *p*” (p. 95). The basic idea of (*), therefore, is that if a luminous condition obtains in a given case, then if one does not already know that the condition obtains, one *could* come to know that it does merely by taking the time to carefully reflect on the matter.

There are many conditions that are uncontroversially non-luminous, such as *that one has hands* (someone who is blind and paralyzed might have hands but not be in a position to know that she does). But there are also many conditions that it is extremely natural to take to be luminous, such as:

- that one is in pain;
- that one feels cold;
- that it appears to one that *p*;

- that one believes that q ;
- that words X and Y have the same meaning for one;
- that one's evidence includes the proposition that r ;
- that one's evidence appears to include the proposition that s ;
- that it is rational for one to believe that t ;
- that it is rational for one to do φ .

Williamson argues that each of these conditions is not luminous; as he sees it, the only conditions that might, perhaps, be luminous are trivial conditions that either obtain in every case (such as the condition *that one exists*) or obtain in none (such as the impossible condition, for which (*) vacuously holds).²

Williamson's central anti-luminosity argument has the following form.³ Let us fix on the condition *that one feels cold*, which Williamson takes to be as good a candidate for a luminous condition as any; later the argument will be generalized so as to apply to almost any putatively luminous condition. Williamson begins by asking us to "consider a morning on which one feels freezing cold at dawn, very slowly warms up, and feels hot by noon" (p. 96). We can stipulate that during this entire process, one does everything one can to decide whether or not one feels cold. We can also suppose that there is no other relevant change in the situation over time: all one does for the entire morning is sit there focusing on how hot or cold one feels as the temperature slowly gets warmer. Moreover, we can make the plausible assumption that one's feelings of hot and cold change so gradually during the

2. Although the details of his position have changed over the years, the philosopher who comes closest to explicitly endorsing a set of luminosity claims that play a central role in his epistemology is Roderick Chisholm; see his (1956), pp. 725–731; (1977), pp. 20–23; (1982), pp. 9–13; and (1997), pp. 27–29. However, most epistemological internalists are also committed to luminosity claims of one form or another.
3. See Williamson (2000), ch. 4. An earlier version of the argument can be found in Williamson (1996).

course of the morning that one is not aware of any change in those feelings from one millisecond to the next. Let $t_0, t_1, t_2, \dots, t_n$ be a series of times at one-millisecond intervals from dawn to noon. For each integer i such that $0 \leq i \leq n$, let α_i be the case at time t_i on the morning in question. Finally, let c be the condition *that one feels cold*, let kc be the condition *that one knows that one feels cold*, and let PKC be the condition *that one is in a position to know that one feels cold*. Williamson's argument then proceeds as follows.

First, Williamson has us assume for *reductio* that the condition *that one feels cold* is indeed luminous. It follows that, for each integer i ($0 \leq i \leq n$), if in α_i one feels cold, then in α_i one is in a position to know that one feels cold:

$$(\text{LUM}) \quad (\forall i, 0 \leq i \leq n)(c \text{ obtains in } \alpha_i \supset \text{PKC obtains in } \alpha_i).$$

Second, Williamson notes that since in each α_i one is doing everything one can to decide whether one feels cold, it follows from his stipulation about the meaning of the expression "being in a position to know" that if in α_i one is in a position to know that one feels cold, then in α_i one does in fact know that one feels cold. Thus we have the following:

$$(\text{POS}) \quad (\forall i, 0 \leq i \leq n)(\text{PKC obtains in } \alpha_i \supset kc \text{ obtains in } \alpha_i).$$

Third, Williamson appeals to what is perhaps the single most important assumption in his entire book. According to Williamson, *knowledge requires safety from error*: in order for one to know something, one must not have been easily wrong in coming to believe it. Much more will be said about Williamson's safety requirement in Sections 3 and 4, below, but for now the following comments will suffice. Suppose that in case α_i one believes that one feels cold. Williamson insists that in order for one's belief in α_i to be safe enough to constitute knowledge, one's belief must not be false in any similar case that one cannot discriminate from α_i . Now one such case is α_{i+1} , the case one millisecond later, so it follows that if one *knows* in case α_i that one feels cold, it must still be *true* in α_{i+1} that one feels cold. As this reasoning will work just

as well for any integer i such that $0 \leq i < n$, the safety requirement on knowledge plus one's limited discriminatory capabilities give us the following *margin-for-error principle*:

(MAR) $(\forall i, 0 \leq i < n)(\text{KC obtains in } \alpha_i \supset \text{c obtains in } \alpha_{i+1})$.

Finally, we have two last assumptions that follow merely from the description of the case:

(BEG) c obtains in α_0 .

(END) c does not obtain in α_n .

That is: at dawn one feels cold, and at noon one does not feel cold.

However, now contradiction looms. By (LUM), if c obtains in α_0 , then PKC obtains in α_0 . By (POS), if PKC obtains in α_0 , then KC obtains in α_0 . By (MAR), if KC obtains in α_0 , then c obtains in α_1 . Therefore from these three conditionals and (BEG), it follows that c obtains in α_1 . Moreover, by a similar chain of reasoning, we may conclude that c obtains in α_2 , that c obtains in α_3 , and so on, until we reach the conclusion that c obtains in α_n . But this contradicts (END). Thus one of the argument's five premises must be false. Williamson claims that (POS), (MAR), (BEG), and (END) are all unassailable, so he infers that (LUM) is the premise responsible for our reaching a contradiction. Conclusion: the condition *that one feels cold* is not luminous — one can feel cold without being in a position to know that one feels cold.

All the above argument assumes about the condition *that one feels cold* is that there exists a continuum of cases, starting from a case in which that condition obtains and ending with a case in which it does not, such that the underlying basis for the condition's obtaining or not obtaining changes so gradually that one cannot discriminate a change in that basis from one case to the next. Thus we can run a parallel argument on any condition for which such a continuum exists. We could argue that the condition *that one is in pain* is not luminous by considering a series of cases in which one feels an agonizing pain that gradually subsides until one feels nothing at all. We could argue that the condition *that it appears that there is a computer in front of one* is not

luminous by considering a series of cases in which one at first clearly sees a computer but then one's eyesight gradually gets blurrier and blurrier. Perhaps we could even argue that the condition *that words X and Y have the same meaning* is not luminous by considering a series of cases in which "two synonyms ... gradually diverge in meaning, as a mere difference in tone grows into a difference in application" (p. 106). And so on. Although the details of how we construct these continua of cases will vary depending upon the condition in question, it seems plausible that if Williamson's argument succeeds in showing that the condition *that one feels cold* is not luminous, analogous arguments could show that almost any other condition is not luminous. In particular, as a process of gradual change can take one from circumstances in which it is rational for one to believe that p to circumstances in which it is irrational for one to believe that p , we can establish that, for any sense of 'rational', the condition *that it is rational for one to believe that p* is not luminous: it can be rational, in that sense, for one to believe a given proposition without one's being in a position to know that it is.⁴

Williamson's anti-luminosity argument forms the backbone of *Knowledge and Its Limits*: not only does he take it to show that no non-trivial condition is luminous, but he also uses the argument (or variants of it) to reply to an important objection to his claim that knowledge is a mental state (ch. 4); to contest Dummett's argument for an anti-realist theory of meaning (ch. 4); to argue against a version of the KK-principle according to which one is always in a position to know when one knows a given proposition (ch. 5); to provide a solution to the surprise-examination paradox (ch. 6); to rebut any argument for skepticism about the external world that assumes that we and the envatted versions of ourselves possess the same evidence (ch. 8); and to buttress his claim that our evidence is all and only what we know (ch. 9). Moreover, similar uses of a safety requirement on knowledge to derive margin-for-error principles play a crucial role in Williamson's theory

4. Williamson insists that such a continuum of cases exists for *it is rational for one to believe that p* at Williamson (2004), p. 315; see also Williamson (2005), pp. 481–482.

of vagueness.⁵ And more generally, Williamson's anti-luminosity argument constitutes a novel way of criticizing a venerable philosophical tradition — a way which Williamson uses to lay the foundation for a radical new theory of (as he puts it) knowledge and its limits.

2. Narrowing the Target

Such is Williamson's core anti-luminosity argument; but is it sound? The argument works by generating a contradiction from the five premises (LUM), (POS), (MAR), (BEG), and (END). The last two of these are undeniable: they follow directly from the set-up of the scenario. (POS), on the other hand, one might doubt: it depends on assuming that in each case one has done everything one can to decide whether one feels cold, and some might insist that doing this takes longer than one millisecond. However, we can easily avoid this worry by extending the amount of time between successive cases while keeping the total number of cases the same (so that the entire process takes longer than a single morning), and once we change the set-up in this way, (POS) seems fine.⁶ Thus the crucial question when evaluating Williamson's argument is whether (MAR) is more plausible than (LUM).

Many people, when first encountering Williamson's argument, find (MAR) independently plausible and so not in need of any further justification. However, resting content with the *prima facie* plausibility of (MAR) reduces Williamson's argument to a mere battle of intuitions over which of (LUM) and (MAR) one finds more plausible, and it would then be open to the defender of luminosity who finds (LUM)

5. It is these margin-for-error principles that allow Williamson to explain why, even though on his epistemicist theory of vagueness each vague predicate has a sharp cut-off point, we can never know where those sharp cut-off points lie. See Williamson (1992a), §5; Williamson (1992b), §6; and Williamson (1994), ch. 8.
6. Actually, matters are slightly more complicated than that, for as long as one's feeling of cold is continually changing between successive cases, we might worry that one will not be coming to a decision about a constant feeling of cold. So if we want to completely avoid this type of worry, we should make one more change to the set-up and, instead of having one's feeling of cold change *continuously* over a given interval, have it change in *incremental steps*.

compelling to simply *tollens* Williamson's *ponens*. Moreover, we should be wary of trusting our intuitions about principles that involve applying vague predicates to very similar cases: after all, the typical sorites premise "One hair can't make the difference between being bald and not bald" is intuitively extremely compelling.⁷ So if the only grounds we had in support of (MAR) were our bare intuitions about its seeming plausibility, those would be very thin grounds indeed, particularly in light of the wide variety of radical consequences that Williamson takes to follow from his anti-luminosity argument.

What we want is an independent motivation for (MAR). As mentioned before, Williamson defends (MAR) by appealing to a safety requirement on knowledge. A recurring theme in the Williamsonian *oeuvre*, from his early work on indiscriminability and vagueness on through his recent material more directly concerned with epistemological matters, is that in cases in which we have imperfect abilities to discriminate between alternatives, a safety requirement on knowledge plus our limited powers of discrimination yield margin-for-error principles.⁸ A margin-for-error principle is any principle that, like (MAR), is of the form "If in case α one knows that p , then in sufficiently similar case β it is true that p ." And the safety requirement is Williamson's way of cashing out the idea that reliability is a necessary condition for knowledge: although in earlier work Williamson talks exclusively of a reliability constraint on knowledge without mentioning the word 'safety', by *Knowledge and Its Limits* he passes freely back and forth between talk of reliability and talk of safety.

What exactly is Williamson's safety requirement? And how exactly does it lead to the margin-for-error principle (MAR)? Unfortunately, this matter is complicated by the fact that Williamson explicates the safety requirement in two different ways, the first involving a coarse-grained, "all or nothing" conception of belief, the second involving a more fine-grained conception in terms of degrees of confidence. So in order to

7. Thanks to Bernhard Nickel for this point.

8. See Williamson (1990), pp. 104–106; (1992a), pp. 223–226; (1994), pp. 226–230; and (2000), pp. 17–18.

assess the case Williamson makes for (MAR), we should consider each of these ways of filling out the safety requirement in turn.

3. Coarse-Grained Safety

Williamson often expresses the idea that knowledge requires safety from error in terms of one's not easily being wrong in similar cases: "if one believes p truly in a case α , one must avoid false belief in cases sufficiently similar to α in order to count as reliable enough to know in α " (p. 100); "in case α one is safe from error in believing that c obtains if and only if there is no case close to α in which one falsely believes that c obtains" (pp. 126–127); "if one knows, one could not easily have been wrong in a similar case. In that sense, one's belief is safely true" (p. 147). Passages such as these strongly suggest that the safety requirement involves the following necessary condition on knowledge:

(C-SAFETY) In case α one knows that p only if, in all sufficiently similar cases in which one believes that p , it is true that p .⁹

When so construed, the safety requirement has an undeniable air of plausibility: if it easily could have been the case that one falsely believes that p — if in an extremely similar way the world might have been, one believes that p and this belief is false — then one's actual belief that p , even if true, does not seem secure enough to constitute knowledge.

Suppose I ask you to think of a number between 1 and 10 and correctly guess that you are thinking of 6; my true belief about which number you are thinking of does not count as knowledge. Why? Plausibly, because there is a very similar case in which I still guess that you are

9. This coarse-grained version of a safety requirement on knowledge is roughly equivalent to one that Ernest Sosa has endorsed in a recent series of articles; see Sosa (1999), among other places. Note, however, that Sosa's formulation of the safety requirement involves ascribing non-standard truth conditions to the subjunctive conditional "S would believe that p only if it were the case that p " — truth conditions that Williamson does not necessarily endorse: cf. Williamson (2000), p. 149.

thinking of 6, but my belief is wrong since you are really thinking of, say, 4. Suppose I decide, out of mere pessimism, that the lottery ticket I just purchased is not a winning ticket; even if I end up being right, my belief that the ticket is not a winner does not count as knowledge. Why? Plausibly, because there is a very similar case in which I still believe out of pessimism that the ticket is not a winner, but the balls determining the winning ticket number bounce slightly differently so as to make my ticket a winner. Of course, how one fixes the similarity relation between cases and what determines the threshold beyond which two cases do not count as sufficiently similar will no doubt be murky matters, but there seems little point in denying that (C-SAFETY) has much to be said in its favor.¹⁰

It is crucial to notice that (C-SAFETY) is not itself a margin-for-error principle. According to a given margin-for-error principle, one does not know that p in some case if, in a certain sufficiently similar case, it is false that p . According to (C-SAFETY), one does not know that p in some case if, in a sufficiently similar case, it is false that p and one believes that p . For all (C-SAFETY) says, one might know that p in some case α despite its being false that p in an extremely similar case α^* , provided that one does not believe that p in α^* . It is only nearby *false belief* that, according to (C-SAFETY), blocks one from having knowledge, not nearby *falsity of what is actually believed*.

Recall that the margin-for-error principle needed for Williamson's anti-luminosity argument to go through is as follows:

10. Several recent articles have attempted to provide counterexamples to (C-SAFETY) as a way of resisting Williamson's anti-luminosity argument. See Brueckner & Fiocco (2002), Neta & Rohrbaugh (2004), Comesaña (2005), and Conee (2005). For Williamson's reply, see §I of his (forthcoming). I agree with Williamson's assessment that the alleged counterexamples offered by these authors are not convincing: each aims to describe two sufficiently similar cases such that one knows a given proposition in the first case and falsely believes that same proposition in the second, but for each pair of cases, either it is far from clear that the first case is a genuine case of knowledge, or it is far from clear that the two cases are sufficiently similar in the relevant respects. (This is all the more true when the safety requirement is modified so that one must have sufficiently similar bases of belief in the two cases: see n. 15.)

(MAR) $(\forall i, 0 \leq i < n)(\text{KC obtains in } \alpha_i \supset \text{c obtains in } \alpha_{i+1})$.

How can one derive (MAR) from (C-SAFETY)? Let us concede to Williamson that each case α_i is sufficiently similar to the case α_{i+1} one millisecond later. Then if BC is the condition *that one believes that one feels cold*, straightforward substitution into (C-SAFETY) yields the following:

(SAF) $(\forall i, 0 \leq i < n)[\text{KC obtains in } \alpha_i \supset (\text{BC obtains in } \alpha_{i+1} \supset \text{c obtains in } \alpha_{i+1})]$.

So in order for us to be able to use the coarse-grained version of the safety requirement to justify the crucial premise of Williamson's anti-luminosity argument, we must somehow get from (SAF) to (MAR).

How to do so is not difficult to see. Williamson's guiding thought is that for cases in which we have limited discriminatory capabilities, the safety requirement on knowledge gives rise to a margin-for-error principle. So what we need is some premise encapsulating our subject's inability to discriminate the cases α_i from one millisecond to the next. The most obvious candidate is as follows:

(BEL) $(\forall i, 0 \leq i < n)(\text{BC obtains in } \alpha_i \supset \text{BC obtains in } \alpha_{i+1})$.

The basic idea behind (BEL) is that because the change from t_i to t_{i+1} in the underlying basis for one's belief that one feels cold is beyond the threshold of one's discriminatory capacities, one's belief at t_{i+1} as to whether one feels cold will be the same as one's belief at t_i . (BEL) is a natural way of articulating the idea that one cannot discriminate case α_i from case α_{i+1} with regards to how cold one feels — that, as Williamson puts it, there is "limited discrimination in the belief-forming process" (p. 127). Moreover, once we have (BEL), (MAR) follows from (SAF), given the additional assumption that knowledge implies belief.

So have we managed to adequately justify (MAR)? No, we have not — for the crucial premise (BEL) is a sorites premise. Indeed, (BEL) by itself is enough to generate a sorites paradox from the undeniable assumptions that (i) one believes that one feels cold in case α_0 (i.e., at dawn), and (ii) one does *not* believe that one feels cold in case α_n (i.e.,

at noon): with n uses of *modus ponens* on an instance of (BEL), we can get from the first of these assumptions to the negation of the second. So we should reject any argument that appeals to (BEL); depending on one's theory of vagueness, the principle is either less than perfectly true or outright false. And if Williamson's anti-luminosity argument does indeed implicitly appeal to (BEL), then that argument is open to the charge of illicitly exploiting the vagueness of the term 'believes', in much the same way as a typical sorites argument illicitly exploits the vagueness of a term such as 'bald' or 'heap'.¹¹

11. A similar objection applies to attempts to save Williamson's argument by appealing to a modalized version of (BEL). It might be insisted that if one believes that one feels cold in case α_i , then even if one does not believe that one feels cold in the case α_{i+1} one millisecond later, there must exist at least one other possible case β similar to α_i in which one has the same qualitative feeling of cold as in α_{i+1} and in which one believes that one feels cold. In symbols:

(BEL') $(\forall i, 0 \leq i < n)(\text{BC obtains in } \alpha_i \supset (\exists \beta \text{ similar to } \alpha_i)[Q(\beta) = Q(\alpha_{i+1}) \wedge \text{BC obtains in } \beta])$,

where $Q(\alpha) = Q(\beta)$ signifies that one's qualitative feeling of cold is the same in cases α and β . (Cf. Williamson [2000], p. 127.) Moreover, it is extremely plausible that whether one feels cold in a given case is determined by one's qualitative feeling of cold, so that if one has the same qualitative feeling of cold in two cases, one feels cold in one of those cases if and only if one feels cold in the other:

(QUAL) $(\forall \alpha, \beta)[Q(\alpha) = Q(\beta) \supset (\text{c obtains in } \alpha \text{ iff c obtains in } \beta)]$.

From (BEL'), (QUAL), (C-SAFETY), and the assumption that knowledge implies belief, we can derive (MAR). However, just as repeatedly iterating (BEL) leads to unacceptable consequences, repeatedly iterating (BEL') leads to unacceptable — or at least highly controversial — consequences. Given the undeniable assumption that one believes that one feels cold in α_0 , repeated applications of (BEL') yields the result that there exists some case β in which one feels as hot as one does in α_n (the case at noon in our original scenario), and yet one nonetheless believes that one feels cold. Indeed, if we tweak our original scenario and stipulate that in the final case things have heated up to such a degree that one's qualitative feeling of hot is the same as it would be were one at the center of the sun, then the defender of (BEL') is saddled with the result that there is 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*. Of course, since the similarity relation is not transitive, β will be very distant from any case in the actual world. But I think we should have serious doubts that such a case is even possible — serious doubts that there could exist a being who counts as having beliefs and experiences, and yet whose beliefs and experiences are as wildly

However, might not Williamson derive (MAR) from (SAF) by means of some principle other than (BEL)? Another obvious candidate is

(KNO) $(\forall i, 0 \leq i < n)(\text{KC obtains in } \alpha_i \supset \text{BC obtains in } \alpha_{i+1})$.

(KNO) and (SAF) together yield (MAR), without our even having to assume that knowledge implies belief. But what is the independent motivation for (KNO), other than its being what Williamson needs in order to derive (MAR)? It's not as if (KNO) encapsulates some dictum about the subject's inability to discriminate between successive cases — or if it does, it only does so derivatively, in virtue of being implied by (BEL) and the claim that knowledge implies belief. And it doesn't seem plausible that (KNO) follows from some more general claim that knowing that p in some case requires that one believe that p in all sufficiently similar cases. Suppose I know a given proposition in the actual case; in a very similar case I might be suspending judgment on the matter, or irrationally clinging to a belief in the proposition's negation. Why should *that* bar me from knowing in the actual case? Though knowledge might indeed require a protective belt of cases in which one does not falsely believe, it is extremely implausible to suppose that, in addition, knowledge requires a protective belt of cases in which one believes.

But more importantly, and independent of the specific failings of (BEL) and (KNO), the more general point is this: if one knows that p in some case, (C-SAFETY) has nothing to say about similar cases in which one does not believe that p ; at some point during the morning one will stop believing that one feels cold; so (C-SAFETY) has nothing to say about whether one really does feel cold from that point on. In particular, as it is incontestable that BC does not obtain in case α_n , (SAF) — and hence (C-SAFETY) — will be completely useless in deriving the conditional “KC obtains in $\alpha_{n-1} \supset \text{C obtains in } \alpha_n$,” which is one of

at odds with one another as they would be in β . To think otherwise is to think that the cognitive and phenomenal realms can come apart from each other to an unacceptable degree. (Compare the discussion of Sosa's thought experiment in Section 5 below.)

the instances of (MAR). The basic purpose of (MAR) in Williamson's anti-luminosity argument is as a bridge principle between cases. From (LUM) and (POS) it only follows that if some condition obtains in a given case, then some other condition obtains *in that same case*; with (MAR), on the other hand, we can deduce that because a certain condition obtains in case α_i , a certain other condition must obtain in the successive case α_{i+1} . However, (C-SAFETY) will be unable to fully undergird (MAR), since (C-SAFETY) can act as a bridge principle between successive cases α_i and α_{i+1} only if one believes that one feels cold in both; as this will not be true for all integers i such that $0 \leq i < n$, we will need some *other* bridge principle to secure (MAR) in those cases, and I claim that whatever this additional principle is, it will be implausible.

One might reply on Williamson's behalf: all we need in order to run the *reductio* is a single case in which one falsely believes that one feels cold, so we don't need to go all the way to α_n . Suppose there is some integer j such that: (i) for every non-negative integer $i < j$, both C and BC obtain in α_i , and (ii) BC obtains in α_j but C does not. Then we could use (BEG), (LUM), (POS), and (SAF) to generate the contradictory result that C both does and does not obtain in α_j . However, who is to say that such a j exists — that as one gradually gets warmer and warmer during the course of the morning while carefully attending to how cold one feels, one stops feeling cold *before* one stops believing that one feels cold? Williamson appears willing to grant to his opponent that there might be “a constitutive connection between the obtaining of the condition [that one feels cold] and one's judging it to obtain” (p. 100), and some candidates for such a constitutive connection rule out the possibility that our subject stops feeling cold before she stops believing that she feels cold.¹² The weakest version of a constitutive connection that has this result is the following:

12. A number of authors in the philosophy of mind literature have recently defended accounts according to which there is a constitutive connection between experience and (the exercise of) certain so-called phenomenal concepts, which in turn gives rise to a corresponding constitutive connection between experience and certain phenomenal beliefs that employ those con-

- (CON) If one has done everything one can to decide whether one feels cold, then one believes that one feels cold only if one feels cold.

Since on the morning in question our subject has done everything she can to decide whether she feels cold, it would follow from (CON) that she never believes that she feels cold on that morning without in fact feeling cold. Of course one might doubt that there exists a constitutive connection of this form between feeling cold and believing that one feels cold, but then we need some independent argument against the possibility of such a connection, which the anti-luminosity argument by itself does not provide. Indeed, as Williamson himself notes (p. 100), typically it is *precisely because* they think that there is a tight connection between certain mental states and beliefs about those states that some philosophers claim the mental states in question to be luminous.¹³ So to simply assume that (CON) is false would beg the question against the defender of luminosity.¹⁴

I conclude that (C-SAFETY) is unable by itself to motivate the margin-for-error principle (MAR). However, at times Williamson talks as if (C-SAFETY) is only a first approximation to the proper coarse-grained version of the safety requirement. "A more elaborate account on such lines," he writes, "would qualify '[one] believes *p*' in the conditional to exclude cases in which [one] believes *p* on a quite different basis

cepts: see Papineau (2002), ch. 4; Chalmers (2003); Block (2006); and Balog (forthcoming).

13. See, for example, Chisholm (1956), p. 725, and (1982), p. 10.

14. In his (2004), Brian Weatherson reaches a similar conclusion that a coarse-grained version of Williamson's anti-luminosity argument would be blocked if on the morning in question one believes that one feels cold only if one feels cold. Weatherson makes the point by considering the possibility that *the very same brain state* might constitute the state of one's feeling cold and the state of one's believing that one feels cold, but we need not make so strong an assumption in order to block the argument: all we need is the much weaker claim (CON). Moreover, Weatherson fails to consider how this insight can be extended to Williamson's official defense of (MAR) in terms of a fine-grained version of the safety requirement.

from the basis on which [one] believes *p* in the case in which [one] putatively knows *p*" (p. 149). Williamson seems to be suggesting that a more accurate version of the coarse-grained safety requirement would read as follows:

- (C-SAFETY') In case α one knows that *p* on basis *b* only if, in any sufficiently similar case α^* in which one believes that *p* on a sufficiently similar basis b^* , it is true that *p*.

Suppose that in some case α one believes that *p* on a given basis, and in a sufficiently similar case β one believes that *p* on a very different basis (perhaps one believes that *p* by perception in case α and by testimony in case β). Then (C-SAFETY') captures the very natural idea that if one's belief is true in α and false in β , one's false belief in β should not impugn the reliability (and hence the status as knowledge) of one's true belief in α : after all, one came to believe that *p* on a very different basis in each case.

Thus the shift from (C-SAFETY) to (C-SAFETY') lends more plausibility to the claim that safety is indeed a necessary condition for knowledge.¹⁵ However, one thing that the shift to (C-SAFETY') does *not* do is provide any help with the criticisms raised earlier against the possibility of deriving (MAR) from a coarse-grained version of the safety requirement. As before, let us grant to Williamson for the sake of argument that (C-SAFETY') is true. Then in order to use that constraint to derive the relevant margin-for-error principle (MAR), we need it to be the case that if in some case α_i on the morning in question one knows that one feels cold on a given basis for belief, then in the case α_{i+1} one millisecond later one believes that one feels cold on a sufficiently similar basis. But if one believes on a sufficiently similar basis that one feels cold, then *a fortiori* one believes *simpliciter* that one feels cold. Thus, just as in our

15. The shift to (C-SAFETY') also provides one more tool in resisting the would-be counterexamples to the safety requirement mentioned in n. 10, since many of those examples fix on two cases in which the bases of one's belief are not sufficiently similar (especially if we individuate bases of belief externally, so that phenomenologically indistinguishable bases of belief need not count as sufficiently similar).

attempt to derive (MAR) using (C-SAFETY) as the only bridge principle between cases, in order to derive (MAR) using (C-SAFETY') as the only bridge principle it must be true that one believes that one feels cold in each successive case. And as this is not true for the morning in question, it will be impossible to use (C-SAFETY') to motivate (MAR) without appealing to some sorites-premise-like principle to the effect that if one believes that one feels cold on a given basis in case α_i , then one believes that one feels cold on a sufficiently similar basis in the subsequent case α_{i+1} . But we should reject any argument that appeals to such a principle. So moving to (C-SAFETY') is of no help in attempting to derive (MAR) from the safety requirement.^{16,17}

16. At one point Williamson mentions in passing a third version of a coarse-grained safety requirement, one that provides even less support for (MAR) than the other two. Williamson writes, "If at time t on basis b one knows p , and at a time t^* close enough t on a basis b^* close enough to b one believes a proposition p^* close enough to p , then p^* should be true" (p. 102). This quotation suggests the following version of a coarse-grained safety requirement:

(C-SAFETY'') In case α one knows that p on basis b only if, in any sufficiently similar case α^* in which one believes a sufficiently similar proposition that p^* on a sufficiently similar basis b^* , it is true that p^* .

(C-SAFETY'') avoids one problem that potentially faces (C-SAFETY) and (C-SAFETY'): we might want to use the safety requirement to rule out lucky guesses about *necessary truths* from being knowledge, but if I correctly guess, say, that 853 is prime, there will be no sufficiently similar cases in which I falsely believe that 853 is prime, since what I believe is true in *every* possible case. (Cf. Sainsbury [1995], p. 595; Weatherston [2004], p. 378; Williamson [2000], pp. 181–182; and Williamson [2005], p. 472.) However, as a means of motivating (MAR), (C-SAFETY'') fares even worse than (C-SAFETY) and (C-SAFETY'). As one gradually feels warmer and warmer on the morning in question, eventually one will stop believing that one feels cold and instead believe a proposition that one might express by saying, "I feel coldish." So (C-SAFETY''), plus various background assumptions about the sufficient similarity of both the propositions believed and the bases on which one believes them, will yield the result that if in the previous case one knew that one felt cold, then in the present case it must be true that one feels coldish. And repeated appeals to (C-SAFETY'') in this way will eventually yield the result that in α_n one feels hot—hardly an absurd conclusion!

17. The recurring problem in these attempts to motivate (MAR) by means of a coarse-grained version of the safety requirement is that, in each case, we need some way of ensuring that one continues to believe that one feels cold from one millisecond to the next on the morning in question. Thus we might ask:

4. Fine-Grained Safety

For most of *Knowledge and Its Limits*, when Williamson discusses the safety requirement it is the coarse-grained version of the safety requirement that is explicitly mentioned. And when Williamson provides a formal model of how luminosity might fail (pp. 127–130), his model invokes a conception of safety formulated in terms of coarse-grained, "all or nothing" belief. However, during Williamson's official defense of (MAR), he instead appeals to a more fine-grained notion of

can we abandon these attempts to derive (MAR) from a coarse-grained safety requirement, and instead attempt to derive it from a safety-like requirement that makes no reference to belief? The safety-like requirement that readily springs to mind is the following:

(SIM) In case α one knows that p only if, in all sufficiently similar cases, it is true that p .

However, (SIM) is extremely implausible, for reasons similar to those given against (KNO): why should its not being the case that p in nearby cases obstruct my knowing that p in the actual case if I don't believe that p , or even engage in any belief-forming process vaguely similar to that which I used in forming or retaining my belief that p , in those nearby cases? Indeed, unless we severely restrict what counts as a sufficiently similar case, (SIM) would appear to block almost any belief from counting as knowledge. For example, suppose I am gazing at a leaf on a nearby tree and truly believe that I am seeing a green leaf. Presumably a case whose only difference from the actual one consists in the leaf in question being brown, or in its not being in front of my gaze since it has already fallen to the ground, is similar in nearly all respects to the actual one; yet if we count such cases as sufficiently similar, we are saddled with the absurd consequence that I do not know that I am seeing a green leaf in the actual case. Of course, we could deem these cases as *not* being sufficiently similar, but then we lose the ability of (SIM) to explain the examples used to motivate a safety-like requirement in the first place—for instance, in the lottery example there are any number of differences between the actual case in which my lottery ticket is a winner and the counterfactual case in which my lottery ticket is a loser that are far greater than a difference in the color or position of one mere leaf. As such we would need some new motivation for taking (SIM) to be a genuine necessary condition on knowledge, and how we provide such a motivation is not readily apparent. But more pressing, the defender of (SIM) who counts the counterfactual leaf cases as not sufficiently similar to the actual case presumably does so partially in virtue of the differences in one's *beliefs* between the actual and counterfactual cases, which would imply that a counterfactual case only counts as sufficiently similar if one has all the same beliefs in that case as in the actual case; however, this makes (SIM) a mere notational variant of (C-SAFETY), and thus all the same problems will arise as before when one attempts to use (SIM) to derive (MAR).

safety: as Williamson puts it, his argument here “depends on applying reliability considerations in a subtler way to degrees of confidence” (p. 127).¹⁸ I suspect that Williamson constructs the official argument in terms of a fine-grained version of the safety requirement specifically because he wants the argument to hold even if something like (CON) is true: the real anti-luminosity argument is supposed to be one that even the staunchest defender of constitutive connections must accept. The problem, however, is that whereas the coarse-grained conception of safety is at least somewhat compelling as a way of articulating the general idea that knowledge requires reliably true belief, the fine-grained version has no such appeal. Or so I shall argue.

It will help if I quote Williamson’s initial justification of (MAR) in terms of degrees of confidence in its entirety (p. 97):

Consider a time t_i between t_0 and t_n , and suppose that at t_i one knows that one feels cold. Thus one is at least reasonably confident that one feels cold, for otherwise one would not know. Moreover, this confidence must be reliably based, for otherwise one would still not *know* that one feels cold. Now at t_{i+1} one is almost equally confident that one feels cold, by the description of the case. So if one does not feel cold at t_{i+1} , then one’s confidence at t_i

18. Williamson has a very peculiar notion of degrees of confidence: although one’s degree of confidence in a given proposition is that which (when there is enough of it) constitutes outright belief in that proposition, Williamson insists that degrees of confidence “should not be equated with subjective probabilities as measured by one’s betting behavior” (p. 98). A better indicator of one’s degree of confidence in a given proposition, he claims, is the degree to which one is willing to use that proposition as a premise in practical reasoning (p. 99). Nothing I say in what follows turns on the distinction between Williamsonian degrees of belief/confidence and degrees of belief/confidence more traditionally construed. Indeed, everything I go on to say will be just as convincing if one replaces all occurrences of the expression ‘degree of confidence’ with ‘degree to which one is willing to use the proposition in question as a premise in practical reasoning’ (although doing so may make it more natural to measure these degrees on a scale from -1 to 1, with 1 representing complete willingness to use the proposition as a practical premise and -1 representing complete willingness to use the proposition’s negation as a practical premise, instead of the 0 to 1 scale that I assume throughout).

that one feels cold is not reliably based, for one’s almost equal confidence on a similar basis a millisecond earlier that one felt cold was mistaken. In picturesque terms, that large portion of one’s confidence at t_i that one still has at t_{i+1} is misplaced. Even if one’s confidence at t_i was just enough to count as belief, while one’s confidence at t_{i+1} falls just short of belief, what constituted that belief at t_i was largely misplaced confidence; the belief fell short of knowledge. One’s confidence at t_i was reliably based in the way required for knowledge only if one feels cold at t_{i+1} .

In this argument, Williamson appears to be implicitly appealing to the following principle:¹⁹

- (F-SAFETY) In case α one’s belief that p with degree of confidence c is reliably based in the way required for knowledge only if, in any sufficiently similar case α^* in which one has an at-most-slightly-lower degree of confidence c^* that p , it is true that p .

Moreover, merely in virtue of the description of the scenario, the following holds for the cases α_i under discussion:

- (CONF) For every integer i ($0 \leq i < n$), if in α_i one has degree of confidence c that one feels cold, then in α_{i+1} one has an at-most-slightly-lower degree of confidence c^* that one feels cold.

(F-SAFETY) and (CONF), together with our usual assumption that each

19. More precisely, Williamson appears to be appealing to a principle according to which one’s belief that p with a certain degree of confidence constitutes knowledge in a given case only if, in any sufficiently similar case in which one believes that p with a sufficiently similar degree of confidence *and on a sufficiently similar basis*, it is true that p . (Compare the discussion of (C-SAFETY) at the end of Section 3.) However, for ease of exposition I shall ignore this complication in what follows, since nothing in my criticism of Williamson’s argument depends on issues concerning the basis of one’s belief.

case α_i is sufficiently similar to its subsequent case α_{i+1} , imply the desired margin-for-error principle, (MAR):

- (MAR) For every integer i ($0 \leq i < n$), if in α_i one knows that one feels cold, then in α_{i+1} one feels cold.

Unlike (BEL) or (KNO), (CONF) seems indisputable, given the description of the situation at hand. However, why should we believe (F-SAFETY)?

The crucial step in Williamson's justification of (F-SAFETY) is his insistence that if in case α one has a degree of confidence that p just barely enough to constitute full-fledged belief, then that belief is not safe/reliable enough to constitute knowledge whenever it is false that p in some sufficiently similar case α^* in which one has a slightly less degree of confidence that p , *even if one's degree of confidence that p in α^* is not enough to constitute full-fledged belief*. It is this feature that allows (F-SAFETY) to be a bridge principle where (C-SAFETY) could not be — that allows us to continue to conclude that one feels cold in successive cases α_i even after one's belief that one feels cold has given out. But why should we withhold the honorific 'reliable' in the kinds of cases Williamson describes? What if one's degree of confidence in its being the case that p perfectly tracks the underlying basis for its being the case that p , so that one's degree of confidence that p falls just short of belief at the precise point at which things fall just short of making it the case that p ? Why would *that* be a situation in which one's initial belief that p is not reliable enough to constitute knowledge? Or slightly more realistically — since Williamson seems to be making the contentious assumption that there is a precise cut-off point above which one's degree of confidence always constitutes full-fledged belief and below which one's degree of confidence always does not constitute full-fledged belief — what if one's belief that p tapers off (as it were) just as its being the case that p tapers off, and in precisely the same way? In such a situation, why should one's lower degree of confidence that p when it is not the case that p in any way impugn the reliability of one's slightly higher degree of confidence that p when it *is* the case

that p ? Of course, how one cashes out this "tapering off" metaphor will depend upon one's theory of vagueness, but the main point remains: (F-SAFETY) deems as unreliable belief-forming mechanisms that appear to be as reliable as they could possibly be.²⁰

To illustrate this point, we can use a slightly altered version of Williamson's own example of a subject gradually feeling warmer on a given morning. The basic idea is that if we are going to go fine-grained with respect to belief, we should also go fine-grained with respect to one's feelings of hot or cold. So let us suppose it were possible to measure the intensity of one's subjective feeling of cold using some set of units — call these units "freezons". To fix on some numbers, let us say that, on the given morning, one's feeling of hot or cold is at a level of 50 freezons at time t_0 (dawn) and a level of -50 freezons at time t_n (noon). Let us also suppose that the following is true: at any time one's degree of confidence that one feels cold on that morning directly correlates to one's subjective feeling of cold as measured in freezons (see Figure 1 on the following page). If we wanted to be precise, we could encapsulate this correlation with the following equation: if $c(t_i)$ is one's degree of confidence at time t_i that one feels cold (measured on a scale from 0 to 1), and $f(t_i)$ one's feelings of hot or cold at t_i as measured in freezons, then $f(t_i) = 100 c(t_i) - 50$ freezons.²¹ Maybe human subjects who are carefully considering how cold they feel would have this sort of a correlation in their freezon/degree-of-confidence levels; maybe they wouldn't. For the purposes of my example it doesn't matter—all I need is for it to be *possible* that there could exist a being with such a cor-

20. Note that in raising these possibilities, I do not mean to be making any claims (at least for now) about the accuracy of any actual human's degrees of confidence over time. Rather, my purpose here is to cast doubt on (F-SAFETY) by describing possible subjects for whom the principle gives implausible results.

21. This assumes that one's degree of confidence that one feels cold is 1 at dawn and 0 at noon, but we could easily adjust the equation so that one's degree of confidence that one feels cold is *nearly* 1 at dawn and *nearly* 0 at noon.

relation between its freezon and degrees-of-confidence levels, which surely is the case.

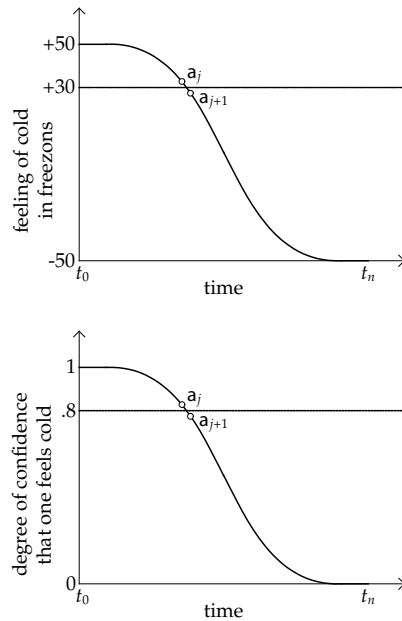


Figure 1

Finally, let us make one last supposition about our scenario. Let one's confidence in how cold one feels that morning be so well-attuned that the following *penumbral connection*²² obtains between the vague expressions 'believes' and 'feels cold': one believes that one feels cold if and only if one feels cold. Then if we sharpen 'believes' and 'feels cold' so that there is a precise cut-off point between the cases in which one does and does not feel cold and a precise cut-off point between

22. See Fine (1975), p. 124. (Note that my invocation of a penumbral connection at this point does not commit me to a supervaluationist theory of vagueness similar to the one defended by Fine in that article: advocates of any theory of vagueness can countenance the existence of penumbral connections.)

the cases in which one does and does not *believe* that one feels cold, the penumbral connection ensures that those cut-off points are the same. Again, it seems evident that one could be a being for whom this is the case.

But now the crucial point comes: if one is such that all of the above is true, then according to (F-SAFETY), at some point during the morning one's belief that one feels cold is too unreliable to constitute knowledge. However, this just seems wrong: one's beliefs about whether one feels cold appear to be as reliable as they possibly could be. This point is most easily illustrated if we sharpen the terms 'believes' and 'feels cold' so that they do, in fact, have sharp cut-off points.²³ Suppose that the sharp cut-off for belief is at 0.8 degrees of confidence: whenever one has a degree of confidence that p greater than 0.8, one counts as believing that p , and whenever one has a degree of confidence that p less than or equal to 0.8, one counts as not believing that p . It follows from our penumbral connection between 'believes' and 'feels cold' that the sharp cut-off for feeling cold is 30 freezons. Now let α_j be the last case during the course of the morning such that, on our sharpening of 'believes', one counts as believing that one feels cold. Then in α_j , one's degree of confidence that one feels cold is $0.8 + \delta$ (for some small real number $\delta > 0$), and one's feeling of cold is at a level of $30 + \varepsilon$ freezons (for some small real number $\varepsilon > 0$). In case α_{j+1} one millisecond later, one's degree of confidence that one feels cold is $0.8 - \delta'$ (for some small real number $\delta' \geq 0$), and one's feeling of cold is at a level of $30 - \varepsilon'$ freezons (for some small real number $\varepsilon' \geq 0$). So (assuming that α_j is sufficiently similar to α_{j+1} and that $0.8 + \delta$ and $0.8 - \delta'$ are

23. Does doing so stack the deck against Williamson? No, it does not. First, as Williamson himself points out (p. 103), if the argument for (MAR) is no longer cogent when we sharpen the relevant vague expressions, then that gives us reason to suspect that (MAR) only seems plausible because it exploits in an illicit manner the vagueness of its key terms (just as the fact that a typical sorites premise becomes obviously false when 'bald' or 'heap' is sharpened reveals that the premise is only plausible because it exploits in an illicit manner the vagueness of 'bald' or 'heap'). And second, on Williamson's own epistemicist theory of vagueness each of the relevant terms already has a sharp cut-off point, so there is no need to sharpen.

sufficiently similar degrees of confidence), Williamson's fine-grained safety requirement (F-SAFETY) implies that one's degree of confidence that one feels cold in α_j is *too unreliable* to constitute knowledge, since in α_{j+1} one's level of cold as measured in freezons slips just below the threshold of what counts as feeling cold (so that it will be *false* in case α_{j+1} that one feels cold). However, this charge of unreliability seems daft: in α_{j+1} one's level of cold as measured in freezons does indeed slip just below the threshold of what counts as feeling cold, but *at precisely the same point* one's degree of confidence that one feels cold slips just below the threshold of what counts as *believing* that one feels cold. Should we then follow Williamson in saying that, "in picturesque terms", the large portion of one's confidence at t_j that one still has at t_{j+1} is misplaced? I think not.

Now one might have legitimate qualms about the notion of freezons that I have helped myself to in this example. How do we measure these freezons? How do we fix their value both inter-personally and intra-personally across time? And what entitles us to move from the familiar fact that some feelings of cold are more intense than others to the more substantive claim that there is a total ordering of feelings of cold by their intensity, let alone a total ordering that has the same structure as an interval on the real line? However, we should note that people can – and have – raised these exact same worries about *degrees of belief*. Freezons are problematic; but so too, in my opinion, are degrees of belief/confidence.

A better objection would be to insist that, once we have sharpened 'believes' and 'feels cold', the ideal epistemic state is not one in which (as in Figure 1) one's degree of confidence that one feels cold directly correlates to one's feeling of cold as measured in freezons, but rather one in which (as in Figure 2) one is absolutely certain that one feels cold until the first case in which one stops feeling cold, at which point one's degree of confidence suddenly drops to 0 and stays that way for the remainder of the morning. But then, the objection continues, since the degree-of-confidence profile given in Figure 1 fails to come close enough to the ideally reliable degree-of-confidence profile given

in Figure 2, it is perfectly acceptable to insist that if one's degree-of-confidence profile were as in Figure 1, one's degree of confidence in case α_j that one feels cold would not, in fact, be reliable enough to constitute knowledge. In this way the counterexample to (F-SAFETY) can be avoided.




Figure 2

It is not clear to me that, even after sharpening the relevant terms, the degree-of-confidence profile given in Figure 2 represents the ideal epistemic state that one might have on that morning.²⁴ However, let

24. Whether this is so depends on issues outside the scope of this paper, such as whether truth is both the aim of *belief* and the aim of *degree of belief*, or whether instead there is something else (degree of truth? objective chance?) that stands to degree of belief/confidence as truth stands to belief.

us grant to the objector that it does; even then we have not managed to save (F-SAFETY).

The objector concedes that if one's degree-of-confidence profile during the course of the morning were as in Figure 2, then in every case in which one's degree of confidence that one feels cold is above the threshold for believing that one feels cold, that degree of confidence would be reliable enough to constitute knowledge. The reason that this possibility is not a problem for Williamson's argument is that, given (CONF), we know that human beings could not have such a degree-of-confidence profile, for that would involve having drastically different degrees of confidence in the same proposition in two successive cases. However, this leaves open the possibility that *other* conceivable beings could have degree-of-confidence profiles that, while not as perfectly accurate as that in Figure 2, are still reliable enough to present a problem for (F-SAFETY). For surely it is not the case that *any deviation whatsoever* from the degree-of-confidence profile given in Figure 2 results in there being at least one case on the given morning in which one believes that one feels with a certain degree of confidence, but that degree of confidence is not reliably enough based to constitute knowledge. To insist upon *that* would be to insist not just that *reliability* is required for knowledge, but moreover that *perfect reliability* is required, and that way skepticism lies. So there must be *some* degree-of-confidence profiles differing from that of Figure 2 such that, for every case in which one has sufficient degree of confidence that one feels cold to count as believing that one feels cold, that degree of confidence is sufficiently reliable to qualify as knowledge.

Another salient feature of the degree-of-confidence profile in Figure 2, other than the fact that it depicts a perfectly reliable belief-forming mechanism, is the fact that the profile is *discontinuous*: at the point at which one goes from feeling cold to not feeling cold, there is a sudden, discontinuous jump from one's having degree of confidence 1 that one feels cold to one's having degree of confidence 0 that one feels cold. Non-idealized physical systems rarely – if ever – exhibit discontinuous phenomena at the macroscopic level, so it seems plausible that,

among the degree-of-confidence profiles for which one's belief that one feels cold always counts as being suitably reliable, some of those degree-of-confidence profiles are continuous. One likely candidate for such a profile is that given in Figure 3: in it, one's degree of confidence over time is nearly as in Figure 2, but the corners of the profile are "rounded off" so as to make the degree-of-confidence function $c(t)$ continuous for all t such that $t_0 < t < t_n$. Note that at this point no assumption is being made that a *normal human* could have such a degree-of-confidence profile; all that is being assumed, at least for now, is that *some* possible creature could.




Figure 3

However, now we have trouble, for the degree-of-confidence profile given in Figure 3 is a counterexample to (F-SAFETY). As before, let α_j be the last of our cases in which one counts as believing that one feels

cold, and let α_{j+1} be the first case in which one does *not* count as believing that one feels cold. Assuming as before a penumbral connection between 'feels cold' and 'believes', α_j will also be the last case in which one feels cold, and α_{j+1} the first case in which one does not feel cold. Moreover, because the degree-of-confidence profile is continuous, we can always choose our α_j and α_{j+1} such that one's degree of confidence that one feels cold in α_{j+1} is slightly lower than one's degree of confidence that one feels cold in α_j ; if this isn't the case when the interval between cases is one millisecond, we can make it so by choosing a smaller interval, such as one microsecond or one nanosecond. Then in α_j one believes that one feels cold with a given degree of confidence, and in sufficiently similar case α_{j+1} one has a slightly lower degree of confidence that one feels cold, despite its no longer being true that one feels cold. So according to (F-SAFETY), one's belief in α_j that one feels cold is not reliable enough to constitute knowledge. However, by assumption the degree-of-confidence profile in Figure 3 is one such that one's belief that one feels cold always meets, for as long as it lasts, the minimum degree of reliability required for knowledge. So, lest we hold that only discontinuous, perfectly reliable degree-of-confidence profiles such as that found in Figure 2 allow one's belief that one feels cold to be reliably based in the way required for knowledge for the entire time it lasts, we must give up (F-SAFETY).

Could one resist this conclusion by insisting that, in order for one's belief that one feels cold to be suitably reliably based for the entire morning, one's degree of confidence that one feels cold must indeed take a discontinuous drop once one stops feeling cold? Such a response is available, but it would divorce reliability talk in the case of knowledge from reliability talk in other domains. A common way of motivating reliability constraints on knowledge is first to note our practice of rating devices (such as thermometers) as reliable when they generally serve their purpose, and then to extend that idea to humans by thinking of our belief-forming mechanisms as nothing more than complex, biological devices for attaining true beliefs and avoiding false ones.²⁵

25. Cf. Williamson (2000), p. 101: "The use of the concept *is reliable* here is a way of drawing attention to an aspect of the case relevant to the application of the

However, by the criterion of reliability being proposed, almost no physical device would count as always being reliable enough for our purposes. For example, if a certain light is built to turn red whenever the temperature in some room is above 32°F, then in order for that light to qualify as sufficiently reliable for the entire time that it is in use, the light would have to discontinuously change in color at the precise instant at which the temperature in the room rises above 32°F; as no physically implemented light could ever do that, no real-world version of such a light would ever count as reliable in this way. But that seems far too demanding a standard of reliability: we can easily imagine a light built for this purpose that we would rely on, and be right to do so, without its color ever making such a discontinuous jump; indeed, we could even build such a light if we wanted.²⁶ And if a requirement of discontinuity is too demanding when ascribing the label 'reliable' to physical devices, why should it fail to be too demanding when ascribing that label to human cognitive systems and their outputs?

Thus I take the various versions of the freezons example to show that (F-SAFETY) does not specify a genuine necessary condition on knowledge. One reply on Williamson's behalf would be to concede the point, but then try to find some alternative to (F-SAFETY) that could still do the work necessary in justifying (MAR). The basic idea would be to argue that, even if Figures 1 or 3 depict the freezon and degree-of-confidence profiles of a creature for whom the condition *that one feels cold* is luminous, humans can never have such profiles, even when they do everything they can to decide whether they feel cold, and so the condition in question is not luminous for creatures like us. However, I have grave doubts about the possibility of such a strategy's

concept *knows*, just as one might use the concept *is reliable* in arguing that a machine *ill* serves its purpose."

26. Williamson might insist that, at the precise instant at which the room's temperature is (say) 31.99°F, our light does not qualify as reliable, though it may well have qualified as reliable several seconds earlier when the room's temperature was (say) 31.5°F. However, why say that? Why not say instead that the light was reliable enough for our purposes the entire time, though of course it could have been a bit more reliable when the temperature was 31.99°F? To repeat a point made earlier: Williamson seems to be conflating *sufficient reliability* at a given time with *perfect reliability* at that time.

ever succeeding. First, notice that it depends upon our being able to find a new criterion for sorting the freezon/degree-of-confidence profiles that result in one's belief that one feels cold being reliable enough for knowledge during its entire duration from those that do not. How to decide upon this new criterion is a delicate affair: it's just not clear what sort of constraints are in play that would allow us to sort the profiles in a suitably precise and yet non-arbitrary manner. But more importantly, once we have this new criterion in hand, we would then need an additional argument establishing that humans always fall on the unreliability side of the divide — that their degree-of-confidence curves over time are *of necessity* too far from the sorts of curves given in Figures 1 or 3 for them to have a belief that they feel cold that is reliable for the entire time it lasts.

It is here that we run into a familiar problem. We saw in Section 3 that no version of Williamson's anti-luminosity argument that attempts to derive (MAR) from a coarse-grained version of the safety requirement would succeed if there is a certain kind of constitutive connection between the coarse-grained states of feeling cold and believing that one feels cold. However, there might also be an analogous constitutive connection *at the fine-grained level* between one's feeling of cold as measured in freezons and one's degree of confidence that one feels cold. And moreover, this constitutive connection might be such that, whenever one has done everything one is in a position to do to decide whether one feels cold, the graphs over time of one's feeling of cold measured in freezons and one's degree of confidence that one feels cold would be as in Figures 1 or 3, or at least close enough to those graphs for one's belief that one feels cold to count as reliable enough for knowledge during its entire duration. Of course, one might doubt that such a constitutive connection exists. But simply to *assume* that it does not, without offering any arguments in support of that assumption, would once again beg the main question at issue, since defenders of luminosity are typically motivated by the thought that there is a tight connection between the obtaining of certain conditions and our beliefs, at least upon reflection, about the obtaining of those

conditions. So even if a suitable replacement for (F-SAFETY) could be found — which itself is highly doubtful — then the brunt of the argumentative work in establishing that conditions such as *that one feels cold* are not luminous would still be left to be done.

5. Do the Relevant Constitutive Connections Obtain?

I have presented two ways of attempting to derive the crucial premise in Williamson's anti-luminosity argument from a safety requirement on knowledge: the first attempted to derive that premise from a coarse-grained safety requirement formulated in terms of "all or nothing" belief; the second attempted to derive it from a fine-grained safety requirement formulated in terms of degrees of confidence. In each case, the original version of the argument ran into problems: in the coarse-grained case, it proved impossible to derive the needed premise (MAR) from the coarse-grained safety requirement without appealing to a dubious sorites-premise-like bridge principle between successive cases; in the fine-grained case, the proffered fine-grained safety requirement failed to specify a genuine necessary condition for knowledge. Moreover, in each case the best attempt at resuscitating a version of the argument — whether in the coarse-grained case by only running the argument until the first case in which it is false that one feels cold, or in the fine-grained case by proposing an alternative to the original (putative) fine-grained safety requirement — was blocked by the possibility that there might be a certain sort of constitutive connection between feeling cold and believing that one feels cold (in the coarse-grained case), or between the degree to which one feels cold and the degree to which one believes that one feels cold (in the fine-grained case).

Settling to what extent, if any, there is a constitutive connection — whether at the coarse-grained or fine-grained level — between feeling cold and believing that one feels cold is beyond the scope of the current paper. But this much seems evident to me: there must be *some* sort of modal connection, constitutive or otherwise, between one's phenomenal experiences (such as one's feeling cold) and one's

cognitive states (such as one's believing that one feels cold) — the phenomenal and the cognitive cannot swing apart from each other any which way one likes. This is brought out by considering a thought experiment once proposed by Ernest Sosa in a very different context. Sosa has us imagine a subject who "has a beautifully coherent and comprehensive set of beliefs", yet through the interference of a group of Cartesian evil demons, has sensory experiences that are "wildly at odds with his beliefs". So, for example, the subject "believes he has a splitting headache, but he has no headache at all; he believes he has a cubical piece of black coal before him, while his visual experience is as if he had a white and round snowball before him," and so on.²⁷ Sosa takes this case to be a counterexample to any theory of justification that makes reference only to one's beliefs and the relations between them, since presumably that theory would deem the subject's beliefs to be fully justified, whereas Sosa insists that there still seems to be something epistemically blameworthy about our subject. However, my reaction to this example is quite different: I fail to see why Sosa's case should trouble anyone, since it seems clear to me that the scenario as described is *impossible* — that there *could not* be a being who counts as having experiences and beliefs even though they are radically disjoint from one another in the way Sosa imagines. Moreover, once one tries to explain why such a being could not exist, one soon finds oneself using phrases like "constitutive connection". Of course, it is one thing to say that a subject cannot have a fully coherent set of beliefs that are wildly at odds with most of her experiences, and quite another to say that, after ideal reflection, a subject cannot believe that she feels cold without really feeling cold. But I take Sosa's example to strongly motivate the idea that there must be some sort of modal connection between the phenomenal and the cognitive realms. Whether that connection is a constitutive one, and whether it is a tight enough connection to block Williamson's argument, are topics for another day.

Note, however, that even if Williamson were somehow able to prove that the relevant constitutive connection does not hold between

the obtaining of the condition *that one feels cold* and one's believing that the condition obtains, in order to extend his anti-luminosity argument to other conditions he would need to argue, on a case by case basis, that an analogous constitutive connection does not exist for each condition to which he applies the argument. For example, in order to use a coarse-grained version of the anti-luminosity argument to show that the condition *it is rational for one to believe that p* is not luminous, Williamson would have to establish that the following is not the case:

- (R-CON) If one has done everything one can to decide whether it is rational for one to believe that *p*, then one believes that it is rational for one to believe that *p* only if it is, in fact, rational for one to believe that *p*.

But for those who have internalist persuasions, the idea that in the limit of inquiry one cannot be wrong about what it is rational for one to believe is extremely plausible. So it will be dialectically difficult, if not impossible, to use a coarse-grained version of the anti-luminosity argument to motivate an externalist conception of rationality according to which it can be rational for one to believe that *p* despite one's not being in a position to know that it is rational for one to so believe. And similar comments apply to attempts to use the anti-luminosity argument (whether in its coarse-grained or fine-grained form) to show that various other conditions, such as *that it appears to one that q* or *that words X and Y have the same meaning for one*, are not luminous: the constitutive connection that Williamson must deny in order to make his argument work is often precisely what is at stake in claiming that the given condition is luminous.

Few would doubt that a version of Williamson's argument can establish that an external-world condition such as *that the temperature outside is less than 90°F* is not luminous. We can imagine a morning on which the outside temperature starts at 70°F and then slowly warms up to 110°F, all while a given subject does nothing but carefully attend to whether the temperature outside feels to be less than 90°F. In this case, the analogue of (MAR) obviously holds: if at any time on the given morning one knows that the temperature outside is less than 90°F,

27. Sosa (1991), pp. 135–136.

then one millisecond later the temperature must still be less than 90°F. Even if somehow one stops believing that the temperature outside is less than 90°F at precisely the point at which the temperature first reaches 90°F, or even if one's degree of confidence in that proposition suddenly takes a sharp drop when the proposition first becomes false, such an occurrence would seem to be nothing more than a fortuitous accident: given the inexactness of our abilities to detect the external temperature around us, there must be some sufficiently similar situations in which one does not stop believing (or in which one's degree of confidence does not significantly drop) at that first point at which the outside temperature stops being less than 90°F. It is only a short step from there to concluding that, regardless of what one's coarse- or fine-grained beliefs are like in the actual world, on the given morning one does not know that the temperature outside is less than 90°F if one millisecond later it is not.

However, for the defender of luminosity, it is not a fortuitous accident that, when one carefully considers the matter, one's beliefs about whether one feels cold perfectly line up with the facts about whether one feels cold. That this could happen would seem miraculous if one holds an *inner perception model of self-knowledge* according to which introspection is fundamentally no different from perception via the five senses: if our faculty for forming beliefs about our own mental states were essentially just an eye turned inward, then given the inevitable inexactness of our external perceptual faculties, it would seem that our inner belief-forming mechanisms must be irremediably inexact as well. But few, if any, of Williamson's targets hold such a conception of self-knowledge.²⁸ Thus for most defenders of luminosity claims there is an important disanalogy between Williamson's argument as applied to the condition *that the temperature outside is less than 90°F* and his argument as applied to a condition such as *that one feels cold*: in the latter case it is *one's own mind* that one is forming beliefs about, and one

28. For a survey of the various views of self-knowledge that do not involve an inner perception model, see Gertler (2003) and the references contained therein.

does not do this by standing outside of oneself, as it were, and using a quasi-perceptual faculty to detect the goings-on within.²⁹

One of the upshots of Williamson's anti-luminosity argument is supposed to be that, because we have no cognitive home of mental states whose obtaining we are always in a position to know of, the type of epistemic access we have to our own mental states is qualitatively no different from the type of epistemic access we have to external-world conditions: the difference is a matter of degree, not of kind. But if my arguments are correct, then there is a sense in which for Williamson's anti-luminosity argument to have any chance of succeeding, he must already assume that conclusion, since he must assume that there is not a special, tight connection between, say, our feelings of cold and our beliefs about those feelings, or between the degree to which we feel cold and the degree to which we believe that we feel cold. If Williamson's anti-luminosity argument succeeds, then it is possible for us to be *epistemically disengaged* (that is, disengaged at the level of knowledge) from our mental states, the meanings of our words, and what rationality demands of us; but he can only establish that conclusion by assuming that, even after ideal reflection, we can be *doxastically disengaged* (that is, disengaged at the level of belief) from a given mental state, the meaning of a given word, or a given demand of rationality.

6. Coda: The Lustrous and the Luminous

I would like to end with a few brief comments about the scope of Williamson's argument.

I take the arguments in Sections 3–5 to show that Williamson does not successfully demonstrate that every non-trivial condition is non-luminous. However, even if I am wrong and Williamson's argument

29. If Williamson's epistemicist theory of vagueness is false, there will be a second important disanalogy between Williamson's argument as applied to the condition *that the temperature outside is less than 90°F* and his argument as applied to a condition such as *that one feels cold*, since the former but not the latter condition will have a sharp cut-off point between the cases in which it obtains and those in which it does not.

does succeed in showing, say, that one can be in pain without being in a position to know that one is in pain, it would not follow that we could use an analogous argument to show that one can be in pain without being in a position to *justifiably believe* that one is in pain. Continuing with Williamson's light-giving metaphor, let us say that a condition *c* is *lustrous* if and only if the following holds:

- (**) For every case α , if in α a condition *c* obtains, then in α one is in a position to justifiably believe that *c* obtains.

A condition that is lustrous shines of its own accord, though not necessarily as brightly as one that is luminous, for the two notions can pull apart: one's position could be such that, if one were to believe that the condition in question obtains, then one's belief would be justified, without the justification that one would thereby have being sufficient to make one's true belief constitute knowledge. And though Williamson's argument might well show that the condition *that one feels cold* is not luminous, a parallel argument cannot be used to show that the condition *that one feels cold* is not lustrous, for while it might be true that knowledge requires safety from error, it is completely implausible to suppose that justified belief requires safety from error.

Or, rather, it is completely implausible to suppose that justified belief requires safety from error, given the common assumption that it is possible to justifiably believe a falsehood. According to the safety requirement on knowledge, one knows that *p* in a given case α only if in some set of sufficiently similar cases, it is true that *p*. What determines that set will vary depending on which version of the safety constraint one is working with: perhaps it is all sufficiently similar cases in which one believes that *p*, or all sufficiently similar cases in which one believes that *p* on a sufficiently similar basis as in α , or all sufficiently similar cases in which one's degree of confidence that *p* is at most slightly less than in α . (See Sections 3–4, above.) But regardless of how one determines the set of sufficiently similar cases, α will itself be among that set, for α itself is as similar as a case can be to

α , and it is trivially true that in α one believes that *p* on the same basis and with the same degree of confidence as one believes that *p* in α . For this reason, the safety requirement for knowledge implies that knowledge is factive: one knows that *p* in some case α only if it is true that *p* in α . And a safety requirement for justified belief would have exactly the same implication: according to such a constraint, one is justified in believing that *p* in a case α only if in some set of sufficiently similar cases, it is true that *p*; as α will be among that set, it will follow that justified belief is factive. However, we can be justified in believing falsehoods: for instance, I might be justified in believing that my copy of *Knowledge and Its Limits* is sitting on my desk at home even if this belief is false because, unbeknownst to me, someone broke into my apartment and stole the book. So there can be no safety constraint for justified belief.³⁰

Thus if we try to show that the condition *that one feels cold* is not lustrous by running an argument parallel to the one that Williamson uses to argue that that condition is not luminous, we will be unable to appeal to a safety requirement for justified belief in order to secure the analogue of premise (MAR), which would read as follows:

- (J-MAR) For every integer *i* ($0 \leq i < n$), if in α_i one is justified in believing that one feels cold, then in α_{i+1} one feels cold.

(J-MAR) simply does not have the same plausibility as (MAR). It seems perfectly plausible that one could *not* feel cold in α_{i+1} and yet have been justified in believing one millisecond earlier that one felt cold; indeed, if α_{i+1} is one of the first few cases on the morning in question in which it is not the case that one feels cold, it seems perfectly plausible that

30. Note that false beliefs can be justified even on Williamson's view according to which (i) evidence is what justifies belief, and (ii) one's evidence is all and only what one knows (" $E = K$ "): even though on this view one's evidence only consists of true propositions, in some circumstances that evidence might be misleading and thus probabilify falsehoods.

one will *still* be justified in believing that one feels cold in α_{i+1} itself and in any number of cases a few milliseconds later.³¹

Therefore a parallel version of the anti-luminosity argument cannot be used to argue that the condition *that one feels cold* is not lustrous — that is, to argue against the claim that whenever one feels cold, one is in a position to justifiably believe that one feels cold. Williamson, however, is likely to be unbothered by this result: after all, one of the chief slogans of his book is “knowledge first” (p. v). For Williamson it is knowledge that is fundamental in all epistemic matters, and interest in justification is only derivative or secondary. However, it is tempting to see Williamson’s anti-luminosity argument as some kind of devastating attack on foundationalism, and it is important to realize that this is simply not the case: most contemporary foundationalists are foundationalists *with respect to justification*, and Williamson’s argument leaves justificatory foundationalism completely untouched. For instance, according to James Pryor’s version of modest foundationalism, “whenever you have an experience as of *p*’s being the case, you thereby have immediate (*prima facie*) justification for believing *p*.”³² And we could easily imagine someone’s defending the related foundationalist view that whenever one has an experience as of *p*’s being the case, one thereby has immediate *prima facie* justification for believing *that one has an experience as of p’s being the case*. Williamson’s argument poses no threat to such claims. When Williamson insists that we are cognitively homeless, what he means is that we are cognitively homeless *with respect to our knowledge of what conditions obtain*: even if his argument works (and it doesn’t: see Sections 3–5), we might still have a cognitive home with respect to matters of justification. I leave it open how satisfactory a cognitive home that would be.³³

31. For this reason appealing to an alleged safety requirement for justified *true* belief will not help the anti-lustrousness argument, either.

32. Pryor (2000), p. 532.

33. For helpful discussion and/or comments, thanks to Mark Budolfson, Alex Byrne, Tyler Doggett, Andy Egan, Alan Hájek, Ned Hall, Patrick Hawley, John Hawthorne, Philipp Keller, Vann McGee, Bernhard Nickel, Dilip Ninan, Agustín Rayo, Daniel Rothschild, Robert Stalnaker, Daniel Stoljar, Eric Swanson, Judith Jarvis Thomson, Timothy Williamson, two anonymous

References

- Balog, Katalin (forthcoming). “The ‘Quotational Account’ of Phenomenal Concepts”. Unpublished manuscript.
- Block, Ned (2006). “Max Black’s Objection to Mind-Body Identity”. In Dean Zimmerman (ed.), *Oxford Studies in Metaphysics*, Vol. 2 (Oxford: Oxford University Press).
- Brueckner, Anthony, and M. Oreste Fiocco (2002). “Williamson’s Anti-Luminosity Argument”. *Philosophical Studies* 110, pp. 285–293.
- Chalmers, David (2003). “The Content and Epistemology of Phenomenal Belief”. In Quentin Smith and Aleksandar Jokic (eds.), *Consciousness: New Philosophical Perspectives* (Oxford: Oxford University Press).
- Chisholm, Roderick M. (1956). “‘Appear,’ ‘Take,’ and ‘Evident’”. *Journal of Philosophy* 53, pp. 722–731.
- Chisholm, Roderick M. (1977). *Theory of Knowledge*, 2nd edition. Englewood Cliffs, NJ: Prentice-Hall.
- Chisholm, Roderick M. (1982). “A Version of Foundationalism”. In *The Foundations of Knowing* (Brighton: Harvester Press), pp. 3–32.
- Chisholm, Roderick M. (1997). “My Philosophical Development”. In Lewis Edwin Hahn (ed.), *The Philosophy of Roderick M. Chisholm* (Chicago: Open Court), pp. 3–41.
- Comesaña, Juan (2005). “Unsafe Knowledge”. *Synthese* 146, pp. 395–404.
- Conee, Earl (2005). “The Comforts of Home”. *Philosophy and Phenomenological Research* 70, pp. 444–451.
- Fine, Kit (1975). “Vagueness, Truth, and Logic”. *Synthese* 30, pp. 265–300. Page references are to the reprint in Rosanna Keefe and Peter Smith (eds.), *Vagueness: A Reader* (Cambridge, MA: MIT Press, 1997), pp. 119–151.
- Gertler, Brie (2003). “Self-Knowledge”. *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/self-knowledge/>

referees for *Philosophers’ Imprint*, and audiences at the 2005 Princeton/Rutgers Graduate Student Philosophy Conference and the 2005 Conference for the Australasian Association of Philosophy. Extra special thanks to Eric Swanson for help in making figures 1–3.

- Neta, Ram, and Guy Rohrbaugh (2004). "Luminosity and the Safety of Knowledge." *Pacific Philosophical Quarterly* 85, pp. 396–406.
- Papineau, David (2002). *Thinking about Consciousness*. Oxford: Oxford University Press.
- Pryor, James (2000). "The Skeptic and the Dogmatist". *Noûs* 34, pp. 517–549.
- Sainsbury, R. M. (1995). "Vagueness, Ignorance, and Margin for Error". *British Journal of Philosophy* 46, pp. 589–601.
- Sosa, Ernest (1991). "Reliabilism and Intellectual Virtue". In *Knowledge in Perspective* (Cambridge: Cambridge University Press), pp. 131–145.
- Sosa, Ernest (1999). "How to Defeat Opposition to Moore". In James Tomberlin (ed.), *Philosophical Perspectives* 13 (Atascadero, CA: Ridgeview), pp. 141–153.
- Weatherson, Brian (2004). "Luminous Margins". *Australian Journal of Philosophy* 82, pp. 373–383.
- Williamson, Timothy (1990). *Identity and Discrimination*. Oxford: Basil Blackwell.
- Williamson, Timothy (1992a). "Inexact Knowledge". *Mind* 101, pp. 217–242.
- Williamson, Timothy (1992b). "Vagueness and Ignorance". *Proceedings of the Aristotelian Society Supplement* 66, pp. 145–162.
- Williamson, Timothy (1994). *Vagueness*. London: Routledge.
- Williamson, Timothy (1996). "Cognitive Homelessness". *Journal of Philosophy* 93, pp. 554–573.
- Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, Timothy (2004). "Replies to Commentators". *Philosophical Books* 45, pp. 313–323.
- Williamson, Timothy (2005). "Replies to Commentators". *Philosophy and Phenomenological Research* 70, pp. 468–491.
- Williamson, Timothy (forthcoming). "Why Epistemology Can't Be Operationalized". To appear in Quentin Smith (ed.), *Epistemology: New Philosophical Essays* (Oxford: Oxford University Press).