

Chapter 4: Anti-Luminosity

An illuminating way of schematizing the whole argument (following Weatherson):

Gradual Change There is a series of cases, each very similar to adjacent cases, that starts with a case where C clearly obtains, and ends with a case where C clearly doesn't obtain.

Luminosity Whenever C obtains you can know it does.

Safety Only safe beliefs count as knowledge, so whenever you can know that C obtains it obtains in all similar cases.

Luminosity and safety imply

Tolerance Whenever C obtains, C obtains in all similar cases.

But tolerance is incompatible with gradual change. Some premise is wrong.

Most of my discussion will focus on three things:

Definition of **luminosity**—

(L) For every case α , if in α C obtains, then in α one is in a position to know that C obtains.

The bridge principle, **I_i**—

(I_i) If in α_i one knows that one feels cold, then in α_{i+1} one feels cold.

What is **reliability**?—

(TWR) Some kind of safety...

Overarching questions for this chapter (non-expository points in bold):

I. Is the definition of *luminosity* that TW uses weak enough and strong enough and correct enough to do the job?

A. Is it weak enough, or could the internalist get by with a weaker claim and swallow that there are no *luminous* states? E.g., coziness (see Hawthorne, and connects with C. and III. below) or other claim in which only some cases of a state are claimed to be accessible (cf. Conee's Severe Pain and Paradigm Pain).

B. Is it strong enough? We want the internalist accessibility claim that is going to turn out to be false not to be true for knowledge of the *external* world. (Connects with II. below because it can turn out not-shown-to-be-false for the external world either for certain notions of reliability.)

C. Is luminosity a useful target? That is, was the claim that the internal world is more accessible than the external world really the claim that with the internal world unlike with the external world we had perfect resolution? (No.) If not then the argument would seem to be irrelevant. However, what then *was* the internalist's claim? (See Hawthorne, Conee.) For one thing the internal seems more accessible because it's "closer in" and "mine". These seem to me different ideas from "home", so they might be useful ways of thinking of how to reformulate internalism.

II. Is the definition of reliability that he uses weak enough and strong enough and correct enough to do the job? Its job is underwriting the bridge principle, I_i .

A. *Safety*. (S's belief in p wouldn't be false in worlds nearby (similar to) the one she believes it in.) α_{i+1} is similar to α_i so if S knows she feels cold in α_i then, by safety, if she believes it in α_{i+1} it must be true. **That she *does* believe it in α_{i+1} doesn't follow from knowledge requiring safety.** Williamson thinks he can help himself to that as an independent part of the description of the case since the change was below the subject's level of resolution. NB: if knowledge requires sensitivity instead, then it won't be true that she believes she feels cold at α_{i+1} if in fact she doesn't feel cold (assuming she knows she feels cold at α_i). So Williamson can't help himself to the assumption that she believes just "by the description of the case," (p. 97) without the further assumption that safety is necessary for knowledge and sensitivity is *not*. **If sensitivity alone is the requirement for knowledge, then TW's argument boils down to the bare fact that we have limited resolution, since that and only that will get you that the subject believes at α_{i+1} (and sensitivity doesn't imply safety). How interesting is the argument then?**

TW argues that safety is better than sensitivity in later chapters, but 1) see my book, and 2) he doesn't seem to think this depends on those arguments. For: TW uses some form of a safety criterion in the anti-luminosity argument but says, in print (pp. 100-101) and in conversation, that there is no need to say more about it. I_i is supposed to be obvious on the basis of intuitions about reliability. (I won't accept that when there's an entire literature on reliability and several different options for what it is. Also, **as we'll see below, he uses safety for one purpose (anti-luminosity), and sensitivity—which is incompatible with fulfilling the first purpose—for a second purpose (relation of reliability to explanation of action).**)

B. *Process reliability*. (S's belief in p is reliable when it was produced by a process that produces true beliefs most of the time.) **Anti-luminosity argument fails if this is what reliability is.** Suppose S's belief that she feels cold is false at α_{i+1} . It doesn't follow that her belief at α_i that she feels cold wasn't process-reliable. Process reliability only requires that the process one used gives correct answers most of the time. So TW will have to go back to the drawing board if this (popular) notion of reliability is favored.

C. *Sensitivity*. S's belief in p is sensitive when if p were false S wouldn't believe it. (Different from safety because counterfactuals don't contrapose.) **I_i fails if sensitivity rather than safety is required for knowledge.** (See A. in this section. But precisely this kind of case might

suggest that sensitivity is too strong a requirement for knowledge.) The bridge principle also fails for external temperature with that requirement, so we wouldn't have here even an argument against luminous *external* conditions. **The reason safety gives Williamson part of the argument and sensitivity doesn't is different directions of fit (with sensitivity we start with and have a requirement on what happens to the subject's belief when p is false at α_{i+1} rather than independently assuming that the subject believes p at that point—if she knew at α_i by sensitivity then she wouldn't believe p if it were false in a similar world).** (See Roush 2005, 121-123.) On the new tracking view sensitivity is probabilistically defined. Does this make a difference? That depends on whether the temperature rising by a millidegree is a *probable* way for it to be false that she is cold. It appears to me that it isn't. Most cases of not being cold will be clear cases rather than borderline cases. On the new tracking view sensitivity is not necessary for knowledge, but the only other way is known implication from something that is tracked, and that isn't the way the subject is being supposed to come to belief at α_i .

III. What is the role of his focus on boundary conditions? Can looking only at such cases get him what he wants?

Observations:

Consider Hawthorne's proposed revision of (L), coziness:

A condition C is *cozy* iff whenever C determinately obtains, you can know that it does.

The internalist view that would be compatible with this would admit that internal conditions are not luminous, but say that we can know (some of) the conditions of our minds whenever they are determinate because some of them are cozy. It's plausible that that's all they wanted anyway. What do you do with an internalist picture of the mind and how we know? You build foundations on it, etc., because it's the firmest knowledge we have. Who thought we had knowledge of indeterminate states or would find it useful?

1. **The luminosity argument appears not to work when we substitute this property for luminosity.** This is even if we give TW that reliability is safety. It doesn't work because there's no lack of resolution from which to conclude that because at α_i the subject knows C obtains, at α_{i+1} the subject believes C obtains because it's so similar she can't resolve. No such thing follows if we look only between determinate states. I.e. the "next" relevant member of the series will be the next determinate case. Determinate cases will not be so similar that she will have to believe at α_{i+1} if she is to be safe at α_i . This points up the fact that in including *all* cases, luminosity is defined as such a strong condition that it may be a straw man.

2. However, I think we can guess how TW is going to respond to this, on the basis of his other views that are not in the book. Namely, he will say that *all* conditions are determinate. (Vagueness is merely epistemic.) So, in particular those ones indistinguishably close to each other are also determinate. **So, coziness is luminosity and there will be determinate states the differences between which are irresolvable.** This is relevant to the section on vagueness in this chapter, where TW argues that vagueness of concepts is not what's responsible for the success of the argument. (To think about: if it

were, why would that matter?) This doesn't directly address, I think, whether the issue of vagueness being real or epistemic is relevant to the argument. TW says in footnote 3 that it isn't, but we've just seen reason to doubt that. It's clear though that he doesn't *want* the anti-luminosity argument to depend on his (rather unpopular) view that vagueness is merely epistemic, a matter of our ignorance rather than something real.

3. However, if TW goes that way, then at some α_i condition C no longer obtains (somewhere, by analogy, midway between bald and not bald). So, what if we had an internalist who was a Williamsonian about vagueness? That would be unfortunate for the internalist because it seems that the luminosity argument would go through (if we gave TW that reliability is safety).

4. However, this reveals that the following may be true:

The luminosity argument works to show there are no luminous states if and only if 1) we grant that reliability is safety, and 2) we assume that vagueness is merely epistemic. This might come from a different argument than that above: notice that when TW argues that "S feels cold" must be true at α_{i+1} he does it by reductio. He assumes that it is false and argues that then S wouldn't know she feels cold at α_i (but she does know by assumption). However, it's not being false doesn't imply it's true if you think vagueness is "real" rather than epistemic. So, I_i doesn't follow. This argument exploits the fact that when vagueness is real truth values aren't guaranteed rather than the fact, as above, that the next determinate state will not be below the subject's level of resolution.

5. Speaking of the reality of determinateness, what about this possibility: that at a certain point, when things get indistinguishable, the subject doesn't have a belief at all about whether she feels cold. The **revised version of internalism would then read**: there are conditions such that if the subject is in a position to believe anything, then she's in a position to know. Just as he assumes that at every point the statement "S feels cold" has a truth value, **TW assumes that at every point S is in a position to have a belief**. So it's an important question whether at α_{i+1} the subject is in a position to have a belief. It seems that she is if we start with the assumption that the difference between the states is unresolvable. However, we can't start there if sensitivity is required for knowledge. But this turns into the same argument as above.

6. On p. 100, TW says all accounts of reliability can ignore the generality problem, basically on grounds that we don't have to be able to analyze a concept in order to use it. However, 1) if all accounts can ignore this problem then we have a question which account is best since (he thinks) they're equal on this dimension, 2) the generality problem for tracking can be solved (see Ch. 3 of my book), and 3) the three popular accounts give different answers as to whether the anti-luminosity argument works.

7. On p. 101, TW argues that a reliability condition on knowledge is consonant with the role of knowledge in the causal explanation of action. "Other things being equal, given rational sensitivity to new evidence, present knowledge makes true belief more likely than true belief does. This is especially clear," he says, "when the future belief is in a different

proposition, that is, when the future belief can differ in truth value from the present belief.” Here reliability is sensitivity, and requiring sensitivity for knowledge is going to give you way more ability to explain action because it means you’ll do the right thing with your belief when the proposition *is* false. Safety doesn’t give you that, for direction of fit reasons. So, TW appeals to reliability as a *sensitivity* requirement on knowledge to account for the ability of knowledge to explain action. However, he appeals to reliability as *safety* to argue against luminosity, and if reliability is sensitivity then the luminosity argument doesn’t work.

8. On p. 102, TW says “A reliability condition on knowledge facilitates the role that knowledge does in fact play in the causal explanation of action. The appeal to such a condition does not depend only on brute intuition; it fits the independently motivated conception of knowing as a mental state.”

I must point out that it fits the denial of the conception of knowing as a mental state just as well. Perhaps even better since TW has had to appeal to two nonequivalent conceptions of reliability, the one of which undermines the purpose of the other in his system. Notice also that these arguments, even if they worked, would only show that TW’s view is possible, not that it’s preferable or plausible.

Conee's Argument Against Luminosity (and $E = K$)

10/04/06 (B.F.)

In "The Comforts of Home" (see website), Earl Conee — an epistemological *internalist* — offers his own, independent argument against luminosity. I think it's worth thinking about Conee's argument and the principles on which it rests, since this gives a nice contrast between Williamson's "externalism" and a contemporary epistemic internalist perspective. Conee endorses the following two (epistemic internalist) principles:

(D₁) If S has evidence against p that is strong enough to defeat whatever grounds S has for believing that p is true, then S is insufficiently justified to know that p is true.

(D₂) If S has a strong unoverridden reason to doubt/reject p , then S does not know p .

(D₁) implies that justification is necessary for knowledge, but (D₂) does not. Conee's arguments can be run using either (D₁) or (D₂), as we'll see shortly. Conee gives the following three ways to defeat the ground for belief provided by awareness of a phenomenal quality:

1. *An argument* "alleging that experience can appear to have a characteristic that is not actually experienced." He describes two possible ways such an argument might run:
 - (a) "A defeating argument can appeal abstractly to experts who are alleged to have discovered a very subtle difference" between "how things seem to be experienced by someone, and how they really are experienced."
 - (b) A defeating argument to the effect "that introspection gives us access to our phenomenal qualities via modes of representation of those qualities, and it has been discovered how to stimulate a brain so as to duplicate such modes in the absence of the quality."
2. "*Sheer testimony*". "One who is known by S to be an expert can quite credibly assert to S that merely apparent phenomenal qualities are possible, without explaining how this can be." And, "there is no conspicuous upper limit on how strong the rational support can be from some such source."

Conee argues that, in such cases, reasons of this kind could defeat the evidence provided by S 's experience of the phenomenal quality (*e.g.*, a chilly feeling) for the proposition (p) that she is experiencing the quality in question (or constitute a strong unoverridden reason to doubt p). Then, by (D₁) [or (D₂)], S would not know herself to experience the quality in question — she would not know p . This is Conee's *internalist* argument against luminosity. Well, almost. He still needs to argue that S may not even be *in a position to know p* in such cases. TW's characterization of "being in a position to know p " is that p be "open to one's view, unhidden, and without any obstacle to one's knowing it." Conee suggests that it is the third of these that fails in the cases he describes. In these cases, S has an obstacle to knowing p : her justification for believing p has been defeated (or she has a strong unoverridden reason to doubt p). And, absent new evidence, Conee thinks this is an obstacle she cannot surmount. Thus, he would argue that these are cases in which S is *not even in a position to know p* . Conee says this argument can be run against the luminosity of any condition.

Moreover, Conee thinks these cases can also be used to show that $E \neq K$. He argues that the thing defeated in these scenarios just *is* S 's "experiential evidence to attribute feeling cold to herself." S still *has* this experiential evidence, but it is defeated (it is, therefore, not a fact that she knows). Thus, S 's total evidence cannot just be identical to the totality of her factual knowledge.