

Appendix D: Weakened Closure and Skeptical Arguments

I here explain how it is, as I claimed in section 22 of Chapter 5, that weakening closure to solve the problem of the aggregation of risk does not undermine the skeptic's argument. It proves convenient to combine that task with responding to a more general, important challenge to the use of closure in skeptical arguments issued by Marian David and Ted Warfield (D&W) in (David and Warfield 2008). D&W argue that when closure principles are weakened in the ways needed to render them plausible, they cease to underwrite classical skeptical arguments of the broad type SSP addresses. This of course constitutes a challenge not only to skeptics, but also to those who, like me, advocate escape routes from skepticism: If the skeptical arguments in question flounder where D&W claim, there is no real threat we need to escape from.

D&W's discussion of various difficulties that come up with formulating closure principles, and how to modify the principles to handle the snags, is extremely helpful. However, they miss what I think is the way for the skeptic to proceed when working with sensibly weakened versions of closure. Accordingly, I will first explain how the skeptic can meet the problems that D&W concentrate on,¹ and then how the skeptic can continue to press on, pretty much undisturbed, after we further weaken closure to handle the problem of the aggregation of risk.

D&W work out their argument most fully in relation to this closure principle:

Closure 6. Necessarily, if S knows that p , and knows that (p entails q), and believes that q based on deduction from p and (p entails q), then S knows that q . (2008: 147)²

If we weaken AI's second premise to make it an instance of this principle, the skeptical argument looks like this:

¹ At (2008: 144-6), D&W do specifically discuss the problem of the aggregation of risk (though not under that label). See n. 12 below for some discussion.

² (David and Warfield 2008: 147). For consistency of presentation, I put D&W's principle into the format we used in Chapter 5 (following Hawthorne), writing out many of the words involved rather than using symbolisms the way D&W do, but I of course follow D&W themselves in matters of content.

Preliminary, invalid version of AI-C6

1. I don't know that I'm not a BIV.

2^{C6}. If

-I know that I have hands, and

-I know that (*I have hands* entails *I'm not a BIV*), and

-I believe that I'm not a BIV based on deduction from *I have hands* and (*I have hands* entails *I am not a BIV*),

then I know that I'm not a BIV.

So, C. I don't know that I have hands.

Our original second premise was just the simple conditional 'bridge' that took us from our base premise, (1), to our conclusion. Weakening (2) to (2^{C6}) renders it too weak a bridge to get us from (1) to (C), leaving the argument invalid.

How shall the skeptic strengthen something in her argument to make up for the weakening of (2) to (2^{C6}), in order to restore validity? It's here that I think D&W miss the best way for the skeptic to proceed. Following what is perhaps the most straightforward strategy, they have the skeptic add these auxiliary premises, asserting that one meets the conditions that have been added to the antecedent of (2^{C6}):

3. I know that (*I have hands* entails *I'm not a BIV*).

4. I believe that I'm not a BIV based on deduction from *I have hands* and (*I have hands* entails *I'm not a BIV*),³

yielding a valid argument, AI-CS (now in its final version), to the conclusion that I don't know that I have hands.

This is indeed a problematic way for the skeptic to proceed. (3) is a safe enough premise, as it's taken to apply to me, or, I suppose, to most readers, recalling that we are understanding the BIV hypothesis so that the entailment holds. *We* know the relevant entailment, after all. (I was able to get away with not addressing this issue in our simplified version of AI because we all realized this.) But as D&W point out, the skeptical argument is typically taken to generalize so as to show that

³ (3) and (4) correspond, respectively, to D&W's (P4) and (P5) at their formulation of the skeptical argument at (2008: 151).

no mere humans know anything like that they have hands (2008: 151), and it seems that there are many people that (3) is not true of. I think the generalizing of the skeptical argument can be fairly well-motivated despite the presence of premises like (3), which don't truthfully apply to all people. However, that's a tricky story,⁴ and, what's more, (4) is in trouble even as it is applied to us, before we get to the generalization move. I mean, I do believe I'm not a BIV, but do even I, or you, really believe this based on the deduction in question? Much here of course depends on the relevant understanding of 'based on', but (4) looks extremely problematic, even as applied to philosophical folks like you and me, actively engaged in thinking about our possible knowledge of such things as that we have hands and that we're not BIVs.⁵

But here is a better way for the skeptic to proceed, which follows the basic way Unger works out an AI-like argument in Chapter 1 of (Unger 1975).⁶ The basic idea: Instead of relying on problematic assertions, like (3) and (4), to the effect that we in fact meet the conditions needed to make closure to go through, the skeptic should instead beef up her base premise so that it claims, not just that one does not know that one is not a BIV, but also that one could not come to know that one is not a BIV through a competent deduction from one's supposed knowledge of the likes of one's hands, even if one met the conditions that need to be added to a sensible closure principle. For closure is supposed to be a necessary truth, so that it covers not just our actual situation, but what would be our situation under various conditions—including, e.g., if we were to perform certain competent deductions

⁴ The skeptic can initially aim her argument at an individual about whom her auxiliary premises are true, reaching the conclusion that this individual does not know that O, and then can try to motivate the thought that if this individual doesn't know, then neither do other people, whether or not the skeptic's auxiliary premises are true of them. This is what D&W call the 'two-stage [skeptical] strategy' that they discuss at 159-163. I take a more positive view than D&W do toward the skeptic's success here, but I agree that this is a tricky call.

⁵ Indeed, as we see in Chapter 7, I think (4) is false of us. Based on how I think the skeptic can best motivate her generalizing move in the second stage of her strategy (see the note immediately above), if one of her auxiliary premises isn't even true of typical philosophically engaged people who are considering the entailment, and so her generalization has to be made from merely hypothetical or actual-but-unusual individuals, it won't be very convincing.

⁶ Recall from note 3 in SSP that in one way, the Argument from Ignorance gets its very name from Unger's presentation of it in that chapter. Unger's use of the basic idea currently under consideration was in large part what my enthusiasm for his presentation was and is based on.

while meeting certain conditions. So a base premise that is strengthened in the way just described could then join forces with a premise based on a plausible closure principle to yield the desired skeptical conclusion.

We can use something like this (named after Unger, who is something of its inspiration, though he is not responsible for any of the details here) as a suitable closure principle:

Closure U. Necessarily, if S knows that p , if p in fact entails q , and if S does not already know that q , then if S knew that (p entails q), and if S were to competently deduce q from p , S would thereby come to know that q .⁷

Yes, there are all sorts of nits to pick about such a principle—and we'll have reason to pick one ourselves shortly (to handle the problem of the aggregation of risk). Its purpose is just to allow us to see how closure principles weakened to respond to the types of considerations that are currently worrying us nonetheless retain their usefulness to the skeptic. Applying Closure U to oneself, while supposing that one is considering whether one is a BIV in light of *I have hands*, and given that *I have hands* does entail *I'm not a BIV*, yields this version of the second premise of the skeptical argument:

2CU. If I know that I have hands, and if I do not already know that I'm not a BIV, then if I knew that (*I have hands* entails *I'm not a BIV*), and if I were to competently deduce *I'm not a BIV* from *I have hands*, I would thereby come to know that I'm not a BIV.

This can combine with a strengthened base premise like

1S. I don't know that I'm not a BIV, and, even if I knew that (*I have hands* entails *I'm not a BIV*), and if I were to competently deduce *I'm not a BIV* from *I have hands*, I would not thereby come to know that I'm not a BIV

⁷ If this principle were only going to be applied to us, who are actively thinking through the issues of whether we know we're not BIVs in light of our presumed knowledge of our hands, it would be more natural to have it read 'if S *knows*'—rather than 'knew'—that (p entails q), and to make other related adjustments. It is phrased as it is to facilitate our later discussion, where we will apply it to people who do not know the entailment, in which application we will want it to say something about what would be the case if they did know the entailment; and also our still later discussion, where we will apply a somewhat modified version of the principle to ourselves, where we will want it to say something about what would be the case if we knew the entailment with perhaps more certainty than we do.

to yield the skeptical conclusion that I don't really know that I have hands. Let us first address the two closely related and pressing questions we face with respect to this argument before we turn to whether the closure principle can be further weakened so as to handle the problem of the accumulation of risk, and yet still be useful to the skeptic.

So, first: Does such an argument generalize so as to plausibly apply to people who are not engaged in thinking about whether they might be BIVs, and who perhaps have never considered such matters? For expository convenience, as in section 6 of Chapter 3, we can just think of this as the issue of whether the skeptical argument plausibly applies to my Mom, now construing her as someone who has never even encountered the BIV hypothesis, and so has no opinions about it.⁸ And our answer is: Yes, we can apply such an argument smoothly to my Mom. Applying Closure U to her, where we do not assume that she has even considered the matter, yields something like:

2M. If my Mom knows that she has hands, and if she does not already know that she is not a BIV, then if she were to consider the matter of whether she's a BIV, and if she knew that (her having hands entails that she is not a BIV), and if she were to competently deduce that she's not a BIV from the fact that she has hands, she would thereby come to know that she's not a BIV,

which could combine with

1M. My Mom does not know that she is not a BIV, and, if she were to consider the matter of whether she's a BIV, and if she knew that (her having hands entails that she is not a BIV), and if she were to competently deduce that she's not a BIV from the fact that she has hands, she would not thereby come to know that she is not a BIV

to yield the skeptical conclusion that my Mom doesn't really know that she has hands. And (1M) seems about as plausible as is (1S).

⁸ Thanks to Michael DePaul, who first pressed me on the issue of whether and how AI-like skeptical arguments can be applied to the non-philosophical, prompting me to first address the issue at (DeRose 2000a: 143-4), in a way that I am updating here. In section 6 of Chapter 3, I address the question of whether anti-skeptical strategies help the non-philosophical to escape from AI-generated skepticism. What DePaul and then, in writing (and more skeptically than DePaul, who perceived the issue, but didn't take a position), D&W, are questioning is whether the nonphilosophical need saving—whether the skeptic's argument works against them in the first place.

This brings us to our second, and most pressing, question, which concerns the tenability of (1S): Does the skeptic's base premise lose its plausibility when it is strengthened from (1) to the likes of (1S) (and (1M))? It may at first seem so. In particular, it can appear that, in strengthening her base premise, the skeptic begins to controversially assume that we cannot get to know that not-H by deducing that not-H from O. But, controversial or not, this assumption is really nothing new. Even when the base premise was the simple (1), taken to apply to you and me, the plausibility of that base premise already depended on our assuming that we couldn't get to know that not-H by deducing that not-H from O. The only reason that we were able to get away with keeping this as an implicit assumption and thereby keeping the argument simple when it was applied to us was that we realized that we had been all along considering whether we could know that we're not BIVs in connection with our presumed knowledge of our hands, and so, if it were possible, would have come to know that not-H by deducing it from O. So, all along, the skeptic has been assuming that folks can't come to know that not-H simply by deducing that not-H from their belief that O. S's grasp of O seems *undermined* by H, as the skeptic might like to put it, for the purpose of knowledgeably fueling an inference to not-H. That's no way to come to know that not-H! Indeed, from some, no doubt unsympathetic, angles, the attempt to come to know that not-H by inference from the likes of O can even look downright ridiculous. The skeptic's assumption, and so (1), enjoy a good deal of intuitive plausibility.⁹ I discuss the intuitive plausibility of (1) in Chapter 2, finding it to be at least very considerable. This plausibility has depended all along on the skeptic's assumption, just identified, and so would be enjoyed as well by a base premise beefed-up in the ways [(1S) and (1M)] we are considering.¹⁰

⁹ For a very different view, see (Roush 2010), where Sherrilyn Roush claims and argues that the skeptic's base premise, as I'm here calling it, loses its plausibility when the skeptical hypothesis is construed, as we are construing it here, in such a way that its falsity really is entailed by the likes of O.

¹⁰ I should note, though, that a consideration I brought forward in Chapter 2 may be somewhat compromised for the purpose of evaluating the plausibility of the strengthened base premises. I was impressed by how many introductory philosophy students judged that they did not know that they weren't BIVs when the issue of whether or not they knew that was presented to them by itself, without reference to any connection, via an AI-like argument, to the issue of whether they knew something like that they had hands. These students, then, may not have been considering the issue of whether they knew they weren't BIVs in connection with their supposed grasp of things like the existence of their hands, rendering it a bit problematic to take their overwhelmingly negative evaluation of whether they knew they were not BIVs to support, not just the likes of

What's more, the sense that one's grasp of *O* is undermined by *H* for the purpose of knowledgeably basing an inference to not-*H*, and the related skeptic-serving intuition that deducing not-*H* from the likes of *O* is no good way to come to know that not-*H*, seems not to be at all based on any lack of security in *S*'s grasp of the entailment (of *not-H* by *O*). This indicates that the skeptic's case won't be hurt when we further weaken closure to handle the problem of the aggregation of risk. For if the skeptic-friendly intuition in question is not at all driven by any thoughts about the insecurity of one's grasp of the entailment, then we would not lose plausibility by strengthening the skeptic's base premise from (1S) to

1S+. I don't know that I'm not a BIV, and, even if I KNEW that (*I have hands* entails *I'm not a BIV*), and if I were to COMPETENTLY deduce *I'm not a BIV* from *I have hands*, I would not thereby come to know that I'm not a BIV,

which is just like (1S) except for having one occurrence of all-caps 'KNEW' where (1S) just has ordinary 'knew', and for having an all-caps instead of ordinary occurrence of 'COMPETENTLY', where 'KNOW' and its cognates designate knowing with absolute certainty, with no risk or micro-risk (or chance or micro-chance or possibility or micro-possibility¹¹) whatsoever of error, and 'COMPETENTLY' designates the kind of absolutely certain deduction (where the deduction of *q* from *p* adds no risk of error at all in one's grasp of *q* beyond what was already present in one's grasp of *p*) befitting of a subject who really KNOWS that *p* entails *q*. This further strengthened base premise can combine with a closure principle that is further weakened to handle the problem of the accumulation of risk. Here we don't need a very precise patch for our closure principle.¹² It's fine for

(1), but the likes of (1S). However, I would have to think that when confronted with even the bare question of whether they know they are not BIVs, while quickly surveying possible ways by which they might possibly know such an exotic thing, folks would very naturally consider whether they might know it through their supposed knowledge of such facts as that they have hands. After all, we are here asking them whether they know some exotic thing, where they have an abundance of things they take themselves to know that entail the exotic item. What could be more natural than to consider whether one's supposed knowledge of those mundane facts might not provide an easy foundation for the exotic thing?

¹¹ See Chapter 5, section 24 for how I use 'micro' here.

¹² A more precise patch might instead add to the antecedent of the (multi-premise) closure principle (letting single-premise closure be just a special instance), not just that *S* knows the premises of the deduction, *p*₁, ..., *p*_{*n*}, and the fact that those premises entail the conclusion, but something like that *S* knows these *n*+1 facts with enough security that if *S* were to deduce the *n*+1-place conjunction of these facts from them, *S* could thereby come to know that conjunction. Alternatively, following a similar principle concerning justification

our current purposes of showing how the skeptic's argument survives the problem for us to 'over-fix' the principle in response to the problem, so long as the over-fixed principle is still useful to the skeptic's cause. Recall that the aggregation of risk is a bigger problem for multi-premise closure than it is for single-premise closure. My own (perhaps controversial) judgment was that even single-premise closure could possibly fail, since some small, knowledge-compatible insecurity in S's grasp of the premise could combine with a similar small insecurity in S's grasp of the entailment, leaving the aggregated risk of being wrong slightly too much to allow for knowledge of the conclusion. But then this could easily be fixed (even if perhaps over-fixed) by stipulating that the entailment must be not only known, but KNOWN, by S, who then not only competently, but COMPETENTLY, deduces her conclusion. This would leave us with something like Closure U-, which I won't display, but is just like Closure U, except for having 'S KNEW that (p entails q)' where U has 'S knew that (p entails q)' and for having 'COMPETENTLY' where U has 'competently'. Applying Closure U- to oneself, while supposing that one is considering whether one is a BIV in light of *I have hands*, and given that *I have hands* does entail *I'm not a BIV*, yields

2CU-. If I know that I have hands, and if I do not already know that I'm not a BIV, then if I KNEW that (*I have hands* entails *I'm not a BIV*), and if I were to COMPETENTLY deduce *I'm not a BIV* from *I have hands*, I would thereby come to know that I'm not a BIV,

which can combine with (1S+) to yield the skeptic's conclusion.

Thus, when closure is weakened in the intuitive way that undermines Hawthorne's argument against our solution to the Harman lottery paradox, it remains useful to the skeptic—which is good news for those of us in the business of providing different escapes from skepticism.

that Feldman considers (Feldman 1995: 448), D&W consider a principle for knowledge, their Closure 3 (David and Warfield 2008: 144), that seeks to avoid our problem by reformulating its antecedent so that it requires the subject to know the conjunction rather than just the individual conjuncts. D&W go on to argue (2008: 145-6) that while such a move makes the skeptic's closure principle more secure, it also makes it more difficult for the skeptic to successfully employ in her argument. But in the text above, I show how a principle weakened to solve (indeed, perhaps to over-solve) our problem can still be successfully utilized by the skeptic.