

Chapter 8 – Skepticism

Williamson is diagnosing skepticism as a consequence of assuming too much knowledge of our mental states. The way this assumption is supposed to make trouble on this topic is that it leads us to think we have the same evidence in the good and bad cases (because evidence is the stuff on the inside we have access to and hence knowledge of, and our evidence is *only* that kind of stuff).

There are dangers of begging the question in what Williamson is trying to do. For example, it's important to see that it just follows automatically from having an epistemologically externalist approach to knowledge that if we *are* in the good case, then we know we're in the good case. This is because all the causal (say) stuff is working right (e.g. perception), and the epistemological externalist makes no further demands in order to count someone as knowing. (Obviously, the epistemological internalist couldn't disagree more with either of these claims.) But no epistemological externalist thinks that *this* is a reply to the skeptic or an argument that the skeptic makes false assumptions. If anything, this looks like a strike against an externalist view to the extent that it shows how impoverished externalism is to explain why the skeptical challenge is appealing at all. (Tracking can, but other externalisms have nothing.)

Williamson is aware of this danger of course, and does not appeal directly to his E=K thesis (so since your knowledge is different your evidence must be different), for example, in order to diagnose and undermine the appeal of skepticism. He's making arguments that look independent of that. But watch carefully for slip-ups, and ask yourself whether the arguments that he says make skepticism appealing, and that he shows don't work, are really necessary for skepticism to be appealing. Also, ask yourself what would an example be of the kind of evidence that is different in the good and bad cases?

Williamson charges that the natural argument the skeptic has that one's evidence is the same in the good case and the bad case is a reduction ad absurdum (pp. 169-70) that assumes that one knows what one's evidence is. In section 8.6 TW waves the sorites wand and reduces this assumption to blubbering foolishness. (Dead horses don't deserve this kind of treatment.)

TW says that this argument is an application of the anti-luminosity argument to sameness of evidence (so see our earlier discussion and criticisms of that). Since I've lost all patience with sorites arguments in this domain, I'll just make a few observations. This argument doesn't seem quite an application, since it goes by constructing an argument that is *parallel* to the argument he attributes to the skeptic. This means we have to ask whether the parallel argument is *relevant* to that sameness of evidence argument ((1)-(12), pp. 171-2). One way that it seems to me it's not relevant can be seen in claim (2_i), p. 175:

(2_i) For any appropriate property p_i , if in a_{i-1} one's evidence lacks p_i , then in a_i one knows that one's evidence lacks p_i .

He says the justification for this claim is just like the justification for claim (2), p. 171:

(2) For any property p_i , if in the good case one's evidence lacks p_i , then in the bad case one knows that in the good case one's evidence lacks p_i .

To say that the argument for (2_i) is the same as the argument for this one is to presuppose that the relevant difference between the good and bad cases--what could tell you which case you're in if only you knew it--is like the difference between light changes at 1 millisecond intervals. This doesn't seem at all right since the difference between the good case and bad case is the difference between having a hand and not having a hand, which isn't a gradual question.

You could say: well, it doesn't matter since all TW needs is one p_i such that people can't know whether p_i is a property of their evidence, in order to undermine the skeptic's sameness of evidence claim, since the sameness of evidence claim depended on the claim that one *always* knows what one's evidence is.

But this doesn't seem to me nearly enough. That people don't know, indeed can't know, these small differences doesn't matter since being able to know them wouldn't help, and because it doesn't imply people can't know big differences, the kind of things that are differences between the good and bad scenarios, if only we could see those differences.

In other words, one doesn't have to know *everything* about one's evidence to know that the good and bad cases are indiscriminable. You don't know properties of the evidence that are indiscriminable from each other, but knowing those wouldn't have helped you tell whether you were in the good or bad case anyway, so we don't need to claim we *do* know them in order to defend an indiscriminability claim. Besides, talk of a "sameness of evidence" claim is a little misleading. All you need to know is that your evidence is the same *as far as you know*.

I'm not sure we needed the anti-luminosity argument against our knowing what our evidence is, since, as TW argues in a later section (8.7), and as has been familiar to Bayesians for some time, rationality is not transparent. That is, one can be rational without knowing that one is. One can be responsive to evidence in a rational way, that is, in conformity with the axioms of probability, for example, without knowing what one's evidence is. (TW is all wrong about Bayesianism in that section.) Thus, we shouldn't have expected in the first place that we know what our evidence is. I'm going to focus on the claim:

(I) For any appropriate property p_i , in any case in which one's evidence has p_i , one knows that one's evidence has p_i .

Roughly, one knows what one's evidence is. In particular, I'm going to ask whether the skeptic needed that claim in order to create the problem. I will argue that far from needing that, he really needs a version of the denial of it, the claim that TW himself is arguing for: one does not in general know what one's evidence is.

There are two different kinds of evidence that TW persistently conflates in these chapters that deal with the concept. One is the notion we invoke when we say the enormous pool of blood is evidence of a murder, more precisely *that there is* a huge pool of blood is evidence *that there was* a murder. In this sense of evidence nothing is implied about whether it is available to anyone. It's this same sense of evidence a detective invokes when she says there must be evidence, if only she could find it. In this sense evidence is facts that are traces or indicators of other putative facts. We may respond rationally and appropriately to such facts without knowing that we do—we do this every day—and without knowing what our evidence, in this sense, is.

If I'm in the good case, then that I have a hand is evidence (in this sense) that I'm not a brain in a vat, that is, that I'm in the good case. If I'm in the bad case, then there is no such evidence, to be sure, since there is no hand. Thus there is *an* asymmetry between the evidence in the good case and bad case, but so far it only applies to the kind of evidence we might call objective, that carries no assumption about its availability to my awareness or about my knowledge *that* I am responding rationally to it, or about my knowledge that it is there.

Nothing in the points I've described says that the evidence in the good case and the bad case could not look the same to me. (Is TW claiming that is metaphysically impossible? I don't think that's the tack.) The significance of this is that we have granted everything Williamson argued for: that there is a notion of evidence that is not phenomenal (yes, this meets TW's demands for the way science treats evidence, in section 8.5), that we do not know what our evidence is in this sense of evidence, and that we do not know what its appropriate properties are.

Yet the skeptic can still make his point. This he can do with the other notion of evidence, which is something like evidence for me, evidence that I'm aware is there, so to speak. We can make the point best, perhaps, by discussing it as what my evidence is *for all I know*. And then the skeptic's point would be to ask how it is I know that what seems to me a hand is evidence that I'm not a brain in a vat. The hand-seeming *is* evidence (in the good case) that I'm not a brain in a vat because in that case all is functioning normally so the hand-seeming is an indicator of my having a hand, and both are indicators of non-*envatment*. In the good case those things *are* evidence. In the bad case the first is not and the second isn't the case. I don't know which they are, so for all I know my evidence is the same.

I should say that I'm not drawing this distinction between two senses of evidence just because it's convenient here. It's also what I endorsed in *TT* and is desirable on other grounds. If you read these chapters of TW with the distinction in mind, you feel that he is sweeping it under the rug, using the term "evidence" indifferently between these two ideas, in order to lull us into giving up reference to a distinction between the metaphysically internal and external. Using words differently is a somewhat okay tactic when you're trying to show people it's possible to see something in a new way, but when you see how he's doing it here, and that the success of his argument actually *depends* on conflating the two notions, then it begins to look like he's massively begging the question.

Notes on Williamson on Evidence and Skepticism

11/01/06 (B.F.)

The central claim Williamson pushes in his scepticism chapter is that there is an *evidential* asymmetry between the good case and the bad case. According to Williamson, an agent in the good case can possess more evidence than an agent in the bad case. For instance, an agent S in the good case can have the evidence that S has hands, whereas, in the bad case (*e.g.*, an envatted S), S cannot have the evidence that S has hands. I think it's difficult to assess Williamson's position in this chapter without assessing his understanding of "The proposition p is part of S 's evidence". We'll do this next week in detail, but even this week I think it is useful to have in mind Williamson's schema:

(EV) E is evidence for H for S iff S 's evidence includes E and E supports H .

What's at issue in the scepticism chapter is just the first part of (EV) — which propositions are "included in S 's evidence". We will assume that the proposition (E) that S has hands *supports* the proposition (H) that S is not in the bad case. What's controversial is whether S can *possess* this evidence in both the good and the bad cases. The sceptic will say S *can* possess E in both cases (if S can possess E at all). Williamson will say S can *only* possess E in the *good* case.

Most of Williamson's moves in the scepticism chapter are defensive in nature. They aim to block certain arguments which aim to show that S 's evidence is the same in the good and bad cases. I tend to agree with Fumerton (and Sherri) about the relative weakness of the defensive moves Williamson offers here. He seems to offer a very naïve reconstruction of the sceptic's reasons for saying that S can have E in the bad cases (if S has E in the good case). *E.g.*,

- (1) For any relevant property P , in any case in which one's evidence has P , one knows that one's evidence has P .

As Fumerton explains, nobody (not even a sceptic) should assume something this strong. And, as Sherri explains, there is a rather strong intuition that there can be pairs of good/bad cases in which the agent "has" the same evidence in what seems to be a salient and legitimate (and interesting) sense of "has".

An analogy might be useful here. Let's say you're someone who values money for its *instrumental* power. You like having money only to the extent that it will actually do things for you. As it turns out, *unbeknownst to you*, you "possess" lots of money. Some benefactor has placed a large sum of money in an account that (legally) you own. But, the catch is that he didn't tell you anything about this, and you will never find out about it.¹ Are you rich? Well, there is a sense in which you are rich, but it's not a very interesting sense, especially by your own lights. *Legally*, you own lots of money. But, since — *as far as you know* — you're the same poor schmuck you were yesterday, this is a sense of

¹We can fill-in the details of a complicated story here if we want to make this plausible: assume it's a joint Swiss bank account you forgot you had, and you were estranged from the benefactor for many years, who was a former business partner of yours with access to the account, *etc.*, *etc.*

“rich” that has no *cash value* for you. I think something analogous is happening in this discussion about what evidence you “have” in the good vs bad cases.

Williamson is quick to point out that we cannot always know what evidence we “have”. And, I (like Sherri) think that the sceptic should simply welcome this and say “Yeah, and that’s what stinks about the sceptical cases. *As far as S knows*, her evidence is just as it would be if she were in the bad case, whether she is in fact in the bad case or not.” The sceptic can grant that *if S* is in the good case, then she “has” evidence that supports she is not in the bad case; whereas, *if S* is in the bad case, then she does not “have” evidence that supports that she is not in the good case. But, it’s the *cash value* of the “having” that’s an issue here. We can similarly grant that *if S* is in the kind-but-forgotten-benefactor (good) case then she “has” money that is sufficient to pay her mortgage; whereas, *if S* is in the same-poor-schumck-she-was-yesterday (bad) case, then she does not “have” money that is sufficient to pay her mortgage. So what? How does that help *S* pay her mortgage — *even if* she happens to be in the good case? It doesn’t. I suspect the sceptic would say something analogous about the extra Williamsonian evidence *S* “has” in the good case. The sceptic is interested in *accessible* differences between bodies of evidence. There may be a sense in which we “have” more evidence in the good case, but the question is whether this is a sense that has any *epistemological cash value*.