

Evidence = Knowledge: Williamson's Solution to Skepticism

Stephen Schiffer
New York University

A single argument template—the *EPH template*—can be used to generate versions of the best known and most challenging skeptical problems. In his brilliantly groundbreaking book *Knowledge and Its Limits*, Timothy Williamson presents a theory of knowledge and evidence which he clearly intends to provide a response to skepticism in its most important forms. After laying out EPH skepticism and reviewing possible ways of responding to it, I show how elements of Williamson's theory motivate a hitherto unexplored way of responding to EPH-generated skeptical arguments. Then I offer reasons to doubt the correctness of Williamson's response.

I. EPH Skepticism

The EPH argument template has as its ingredients an uncontentious fact E, a run-of-the-mill proposition P, and a skeptical hypothesis H such that:

- common sense supposes that one would know and be justified in believing P on the basis of E;
- H entails both E and not P; and
- it appears that if E can't justify one in believing not H, then there is nothing else available to justify one in believing not H.

Using these ingredients, the skeptic argues as follows, where by stipulation the subject "I" is a rational thinker who is fully and actively aware that P and H are incompatible and that H entails E:

- (1) I'm not justified in believing P unless I'm justified in believing not H.

- (2) I'm not justified in believing not H unless something other than E justifies me in believing not H.
- (3) There is nothing other than E to justify me in believing not H.
- (4) ∴ I'm not justified in believing P.
- (5) If I'm not justified in believing P, then I don't know P.
- (6) ∴ I don't know P.

When we take E to be the proposition that I'm having such-and-such sensory experiences as of a red cube, P the proposition that there is a red cube before me, and H the hypothesis—call it BIV—that I'm a brain floating in a cubeless vat of nutrients and attached to a device that is causing me to have the such-and-such sensory experiences as of a red cube which I'm now having, then we get:

External World

- (1) I'm not justified in believing that there's a red cube before me unless I'm justified in believing not BIV.
- (2) I'm not justified in believing not BIV unless something other than the fact that I'm currently having such-and-such sensory experiences as of a red cube justifies me in believing not BIV.
- (3) There is nothing else to justify me in believing not BIV.
- (4) ∴ I'm not justified in believing that there is a red cube before me.
- (5) If I'm not justified in believing that there is a red cube before me, then I don't know that there is a red cube before me.
- (6) ∴ I don't know that there is a red cube before me.

By a straightforward extrapolation and generalization, the argument may be continued to show that no facts about my sensory experiences can justify me in believing anything about the external world, and that therefore my sensory experience can't give me knowledge of the external world. And if "I"—that is, a rational thinker in epistemically

optimal conditions—can't have knowledge or justified belief of the kind in question, then no one can have such knowledge or justified belief.

When we take E to be the fact that I have egg on my shirt and seem to remember eating eggs for breakfast, P the proposition that I had eggs for breakfast, and H the hypothesis—call it NEW—that the universe just this moment came into existence, completely as is, with the fact that I have egg on my shirt and seem to remember eating eggs for breakfast, then we get:

Past

- (1) I'm not justified in believing that I had eggs for breakfast unless I'm justified in believing not NEW.
- (2) I'm not justified in believing not NEW unless something other than the fact that I have egg on my shirt and seem to remember eating eggs for breakfast justifies me in believing not NEW.
- (3) There is nothing else to justify me in believing not NEW.
- (4) ∴ I'm not justified in believing that I had eggs for breakfast.
- (5) If I'm not justified in believing that I had eggs for breakfast, then I don't know that I had eggs for breakfast.
- (6) ∴ I don't know that I had eggs for breakfast.

By a straightforward extrapolation and generalization, the argument may be continued to show that no facts about the present can justify me in believing anything about the past, and that therefore such facts can't give me knowledge of the past. And if I—a rational thinker who is actively aware of relevant entailments—can't have knowledge or justified belief of the kind in question, then no one can have such knowledge or justified belief.

When we take E to be the fact that Al broke his toe by stubbing it against a rock and is screaming and writhing on the ground, P the proposition that Al is in pain, and H the hypothesis—call it ZOMBIE—that Al is a zombie who has no sentient mental states, even though he broke his toe by stubbing it against a rock and is screaming and writhing on the ground and in general behaves in ways I expect sentient humans to behave, then we get:

Other Minds

- (1) I'm not justified in believing that Al is in pain unless I'm justified in believing not ZOMBIE.
- (2) I'm not justified in believing not ZOMBIE unless something other than the fact that Al broke his toe and is screaming and writhing on the ground justifies me in believing not ZOMBIE.
- (3) There is nothing else to justify me in believing not ZOMBIE.
- (4) ∴ I'm not justified in believing that Al is in pain.
- (5) If I'm not justified in believing that Al is in pain, then I don't know that Al is in pain.
- (6) ∴ I don't know that Al is in pain.

By a straightforward extrapolation and generalization, the argument may be continued to show that no facts about the behavior of, or causes acting on, another body can justify me in believing anything about the sentient mental states of others, not even that others have such states, and that therefore such facts can't give me knowledge of other minds, not even that there are other minds. And if "I" can't have knowledge or justified belief of the kind in question, then no one can have such knowledge or justified belief.

When we take E to be the fact that all observed ravens have been black, P the proposition that the next observed raven will be black, and H the hypothesis—call it NONUNIFORMITY—that, while all observed ravens have been black, no observed ravens after now will be black, then we can even get:

Induction

- (1) I'm not justified in believing that the next observed raven will be black unless I'm justified in believing not NONUNIFORMITY.
- (2) I'm not justified in believing not NONUNIFORMITY unless something other than the fact that all observed ravens have been black justifies me in believing not NONUNIFORMITY.

- (3) There is nothing else to justify me in believing not NONUNIFORMITY.
- (4) ∴ I'm not justified in believing that the next observed raven will be black.
- (5) If I'm not justified in believing that the next observed raven will be black, then I don't know that the next observed raven will be black.
- (6) ∴ I don't know that the next observed raven will be black.

By a straightforward extrapolation and generalization, the argument may be continued to show that no facts about past regularities can justify me in believing that any past regularities will continue to hold, and that therefore such facts can give me no knowledge of the future. And if "I" can't have knowledge or justified belief of the kind in question, then no one can have such knowledge or justified belief.

It's important to appreciate that the different skeptical paradoxes are instances of the same argument form, because that motivates a defeasible expectation that if any one of the four skeptical arguments goes wrong in a particular way, then they all go wrong in that way. In other words, we should not expect there to be one solution to the problem of the external world and a different solution to, say, the problem of other minds. Skeptical arguments may take other forms, but it's reasonable to suppose that a resolution of the problems raised by EPH arguments will have application to the skeptical arguments that take those other forms, and that any fully adequate response to those other arguments will have application to the EPH arguments.

For the rest of this paper I shall focus just on the EPH argument template as it concerns justified belief, that is, on the template:

EPH

- (1) I'm not justified in believing P unless I'm justified in believing not H.
- (2) I'm not justified in believing not H unless something other than E justifies me in believing not H.

- (3) There is nothing other than E to justify me in believing not H.
- (4) ∴ I'm not justified in believing P.

It's not that no one has thought to challenge premise (5) (If I'm not justified in believing P, then I don't know P). Some have challenged the claim that knowing P entails believing P, while others have conceded that knowing P entails believing P but have challenged the claim that knowing P entails being *justified* in believing P. I'm not sure why no one has thought to challenge the claim that knowing P entails that P is true, since we sometimes say such things as 'I *knew* she would say yes' when we know that she didn't say yes, and that sort of use of 'know' is pretty much on all fours with the examples that are supposed to loosen the ties between knowledge and belief or justification. In any case, I'm not aware of any good reasons to deny (5); Williamson evidently wouldn't deny it; and even if (5) were false, an argument that concludes that we can't be justified in believing the things we're certain we are justified in believing is itself, needless to say, a skeptical paradox worth grappling with.

EPH is a valid argument form, and we come to the skeptical arguments already believing that their conclusions are false. The plausibility of the premises, however, must be earned. We will be in a better position to appreciate the nature of Williamson's innovative response to EPH skepticism if I first briefly review enough of what might be said in support of each of the three EPH premises to justify treating the EPH arguments as *paradoxes*—valid arguments with apparently true premises and apparently false conclusions (thereby showing, once again, that you can't always go by appearances).

Premise (1) (I'm not justified in believing P unless I'm justified in believing not H). Recall that the "I" of the argument is by stipulation a rational thinker who is actively aware that P entails not H. Given that, the plausibility of (1) is entailed by the plausibility of the closure principle:

JBC For any propositions P, Q, one who is actively aware that P entails Q is justified in believing P only if she is also justified in believing Q.

It's easy to see why JBC is plausible. If one may be justified in believing P but not Q when one is actively aware that P entails Q, then it needn't be irrational for one to believe

that P is true and to doubt whether Q is true even while being fully and actively aware that it's impossible for P to be true unless Q is true. But it's doubtful that such a combination of attitudes is possible, let alone can be rationally held.

I do need to say something about the intended meaning of 'is justified in believing' as it occurs in EPH, and thus in JBC. Three justification relations need to be distinguished:

IS E *is* a justification for S to believe P

HAS E is a justification that S *has* to believe P

IN E justifies S *in* believing P

On the intended reading of IS, E can be a justification for S to believe P even though S isn't aware of E and doesn't believe P. For example, a certain symptom may be conclusive evidence that I have a certain disease, even though I'm unaware of the symptom, unaware that I have the disease, and would be unaware that the symptom was evidence of the disease even if I were aware of it.

On the intended reading of HAS, in order for E to be a justification that S *has* to believe P, S must know, or at least believe, E—or at least simply have E, if E is an experience or sensation. (If we assume Williamson's theory, wherein only evidence can justify and a proposition belongs to S's total evidence just in case S knows it,¹ then E is a

¹ Williamson's equating a person's evidence with her knowledge—an equation he calls $E = K$ (185)—is a cornerstone of his theory of evidence (unless otherwise noted, all page references for Williamson are to *Knowledge and Its Limits*). The doctrine that only evidence can justify belief occurs in a few places (e.g. 208), and is evidently also pretty central to Williamson's theory. It has, however, a weak and a strong reading. The weak reading is that if E justifies one's believing P, then E is known, and thus, by $E = K$, belongs to one's total evidence. The stronger reading is that if E justifies one's believing P, then E is evidence for P for one. Williamson sometimes gives the impression that he accepts the stronger reading, as when he says that "evidence for a mathematical conjecture may consist of mathematical knowledge" (207), but here he is probably using "evidence" in a loose vernacular way, since on his account of evidence, nothing can be evidence for a mathematical proposition, and no mathematical proposition can be evidence for any proposition. This is because E is evidence for P only if it raises the probability of P, in the sense that $\text{Prob}(P/E) > \text{Prob} P$, and for Williamson every mathematical proposition has probability 1 or 0. The fact that nothing can be evidence for a mathematical proposition may be taken to be a problem for Williamson's theory of evidence, since we may well want to say such things as that my evidence for P's being a theorem of number theory is that the brilliant number theorist Jones told me that it was.

justification that S *has* for believing P iff E is a justification for S to believe P and S knows E.)

On the intended reading of IN, E can be a justification that S has for believing P and yet not be what justifies S *in* believing P, even though S is justified in believing P. For example, S may know a certain fact about a DNA fingerprint found in a hair sample at the scene of a murder; this fact may be virtually conclusive evidence that the chauffer was the murderer; and S may not know that the DNA fingerprint implicates the chauffer yet still be justified in believing that the chauffer committed the murder on the basis of knowing that two independent and uninvolved witnesses say they saw the chauffer commit the murder. To be *justified in* believing P is to believe P and to be justified in doing so, and for E to be what justifies S in believing P it must be that S believes P on the basis of S's having, knowing, or at least believing, E, in a sense of 'on the basis of' that awaits explication but can be used to sort cases. Assuming Williamson's theory of evidence (see below), we may say that evidence E justifies S in believing P only if S's knowing E accounts for the fact that S is justified in believing P. This in turn implies that, all other things being equal, if E justifies S in believing P, then S wouldn't be so justified if S didn't know E.

Premise (2) (I'm not justified in believing not H unless something other than E justifies me in believing not H). The argument for premise (2) is this:

- (i) We may take it as given that (a) I know for certain that H entails E; (b) I come to know E at a certain time t^{*2} ; and (c) prior to t^* , both E and H were uncertain to me.
- (ii) If (i), then E is evidence *for* H for me at t^* .
- (iii) E can't be part of what justifies me in believing *not* H at t^* if E is evidence *for* H for me at t^* .

We needn't, however, bother about any of this, since all the issues in this paper about what a person is justified in believing pertain only to contingent propositions.

² I hope my use of 'a certain time t^* ' is clear enough. To do the quantification over times properly would make for a less neat statement of the argument. I shall omit temporal references when doing so is harmless.

- (iv) \therefore I'm not justified in believing not H at t^* unless something other than E justifies me in believing not H at t^* .

Only (ii) and (iii) need justification. Let t^* continue to be the time alluded to in the argument, and let

Prob_{old} = probability on all the evidence acquired up to the time just before t^* , the time at which $\text{Prob}(E)$ becomes 1.

Now, it's a theorem of probability theory that

$$[(\text{Prob}_{\text{old}}(E/H) = 1) \ \& \ (0 < \text{Prob}_{\text{old}}(E) < 1) \ \& \ (0 < \text{Prob}_{\text{old}}(H) < 1)] \rightarrow \text{Prob}_{\text{old}}(H/E) > \text{Prob}_{\text{old}}(H),$$

and it's transparently plausible that

E is evidence for H for S at t^* if (a) S knows E at t^* and (b) $\text{Prob}_{\text{old}}(H/E) > \text{Prob}_{\text{old}}(H)$,

and from those two things (ii) follows. Two points also secure (iii): First, E is evidence *against* not H for me if E is evidence *for* H for me (this is reflected in probabilistic terms by the theorem that $\text{Prob}(H) + \text{Prob}(\neg H) = 1$). And second, E can't be part of what justifies me in believing not H if E is evidence *against* not H for me—after all, if E is evidence against not H for me, then E should *lower* my confidence in not H, and something that lowers my confidence in a hypothesis can't be part of what justifies me in believing it.

Premise (3) (There is nothing other than E to justify me in believing not H). One can't provide a prima facie justification for the instances of premise (3) in question without regard to the particular values of H and E. But in each case there is the same pattern of argument: There is nothing other than E to justify me in believing not H because (a) if there's to be evidence against H, it will ultimately come down to evidence of the kind to which E belongs for propositions of the kind to which P belongs; but (b) as the instance of EPH shows, no E-type fact can justify one in believing a P-type proposition unless there is a justification for disbelieving an H-type hypothesis that is independent of that E-type fact. Here's how this plays out with respect to External World:

- (1) Since BIV is a contingent empirical hypothesis, I would be justified in disbelieving it only if I have empirical evidence against it.
- (2) But any such evidence would itself have to consist in propositions belief in which was directly or indirectly justified by my sensory experience.
- (3) Since any such experience will encounter its own BIV hypothesis, if any sensory experience could justify me in believing not BIV, my such-and-such sensory experiences as of Cube could, too.
- (4) But, as we've seen, my such-and-such sensory experiences as of Cube can't justify me in believing not BIV.
- (5) So, there is nothing other than my such-and-such sensory experiences as of Cube to justify me in believing not BIV.

II. Responses to EPH Skepticism

EPH skeptical arguments presuppose that whatever justifies you in believing P if you really do know P will also justify you in believing P if the skeptical hypothesis H is true, and vice versa. Let's call this the *same-justification assumption* (SJA). It will be helpful to restate SJA in the following Williamson-inspired way. In order to make my discussion more concrete, I'll be concerned with SJA only with respect to External World. The EPH skeptic launches External World against the presupposed background of a *best-case scenario* (BCS): a scenario that is consistent with BIV but is otherwise optimal for my having perceptual knowledge that there is a red cube before me, if there is a red cube before me (so if there is a red cube before me and I can't know it there, then I can't know it anywhere). The skeptic then envisages two incompatible further descriptions of BCS. Let's adopt Williamson's untendentious labels and call one of these further descriptions Good and the other Bad:

Good: BIV is false and I perceive, and thus know, that there is a red cube before me [Cube, for short].

Bad: BIV is true (and therefore I don't know Cube), but otherwise my situation is as much like Good as it's possible for it to be.

Then SJA, applied to BCS, has it that, if anything justifies me in believing Cube in either realization of BCS—that is, either in Good or in Bad—then it also justifies me in believing Cube in the other, and that that one thing is the fact that I'm having such-and-such sensory experiences as of Cube.³

Most responses to External World accept SJA.

This is true of the contextualist response, which finds indexicality in External World and holds that it expresses a sound argument in certain “high-standards” contexts, such as a context in which skepticism is being discussed (like there's another context in which you'd find External World?).

It's true of the only way of denying premise (1)—viz., denying the closure of justified belief under known entailment—which I think is advocated only by Fred Dretske.⁴

It's true of the only way of denying premise (2), the “dogmatism” whose proponents include Tyler Burge, Christopher Peacocke, John Pollock, and Jim Pryor.⁵

It's true of the inference-to-the-best-explanation response to premise (3), which claims that, while the evidence is the same in Good and Bad, we're justified in believing not BIV because the commonsense explanation of the fact that I'm having my sensory experiences as of Cube, according to which my experience is a veridical perception, can be said to be a better explanation than BIV of that fact just on the basis of an invidious

³ It seems clear that if one can be justified in believing Cube in Good, then one can also be justified in believing Cube in Bad, *if one can entertain* Cube in Bad. A person who only recently became disembodied and envatted can entertain Cube, but it may be arguable that a creature who has been envatted its entire life wouldn't have the concepts required to entertain Cube. This issue matters with respect to a possible problem for Williamson that I discuss below, but otherwise we should understand BIV in a way that allows for me to entertain Cube in Bad.

⁴ Dretske (1970). Robert Nozick (1981) denies closure for knowledge but not for justified belief.

⁵ Burge (2003), Peacocke (2004), Pollock (1986), and Pryor (2000).

comparison of BIV with the commonsense explanation with respect to theoretical virtues and vices.⁶

And it's true of the response to premise (3)—flirted with by Crispin Wright and accepted by others⁷—which holds that by default we're *a priori* justified in disbelieving BIV simply by virtue of the presuppositional status the commonsense material world hypothesis enjoys in our belief system.

For the record, I don't find any of these responses plausible.

There are, however, in principle two different ways to deny SJA, and either would enable one to challenge premise (3). One way is to claim that I have *no* justification for believing Cube in Good which I also have in Bad, even on the assumption that the fact that I'm having such-and-such sensory experiences as of Cube justifies me in believing Cube in Bad. This is the so-called *disjunctivism* of John McDowell and others, according to which I may be justified in believing Cube, and thus in disbelieving BIV, in both Good and Bad, but that what justifies me is different in the two cases, a perception of a red cube in the one case, and a hallucination of a red cube in the other, these not being states that share a qualitative sensory state that would itself provide justification for believing Cube.⁸ I think that disjunctivism is an implausible response to skepticism for several reasons, two of them being that it yields no response to Induction and an extremely strained response to Other Minds.

The second way of denying SJA holds (i) that even in Bad I am justified in believing Cube, (ii) that what justifies me in believing Cube in Bad is the fact that I'm having such-and-such sensory experiences as of Cube, (iii) that that justification is also a justification for believing Cube which I have in Good, but (iv) that in Good there is another, knowledge-securing justification for believing Cube, which I don't have in Bad, that justifies me in believing Cube. No one to my knowledge has ever responded to EPH skepticism in this way—*unless* this is the response to which Timothy Williamson is committed by the theory of evidence and the reply to skepticism he advances in *Knowledge and Its Limits*.

⁶ See, e.g., Vogel (1990).

⁷ Wright (2004), and see, e.g., White (forthcoming).

⁸ McDowell (1998); see also Martin (2004).

III. $E = K$ and Williamson's Implied Response to EPH Skepticism

With respect to the skeptical argument External World, the *same-justification assumption* (SJA) holds that whatever *justifies me in believing* Cube if I really do know Cube will also justify me in believing Cube if the skeptical hypothesis H is true, and vice versa.

Williamson doesn't consider skepticism in its EPH form, nor does he explicitly consider SJA. He does, however, consider a version of external world skepticism which accepts what we may call the *same-evidence assumption* (SEA)—viz., that the *evidence* one has in Good is exactly the same as the evidence one has in Bad. Applied to External World, SEA holds that whatever is evidence for Cube for me in Good is also evidence for Cube for me in Bad, and vice versa.

In responding to the version of skepticism he considers, Williamson argues that the skeptic goes wrong in accepting SEA. The falsity of SEA follows from two tenets of Williamson's theory of evidence:

EV: E is evidence for P for S iff (i) S's evidence includes E and (ii) $\text{Prob}_{\text{old}}(P/E) > \text{Prob}_{\text{old}}(P)$.

$E = K$: S's evidence includes E iff S knows E.

For suppose I know Cube (and am thus in Good and not in Bad). Then, by $E = K$, my evidence includes Cube. And since it's clear that $\text{Prob}_{\text{old}}(\text{Cube}/\text{Cube}) > \text{Prob}_{\text{old}}(\text{Cube})$, it follows from EV that Cube is evidence for Cube for me. Indeed, I can have no better evidence for Cube than Cube: since $\text{Prob}_{\text{new}}(\text{Cube}) = \text{Prob}_{\text{old}}(\text{Cube}/\text{Cube}) = 1$, Cube is conclusive evidence for me that Cube. Thus, I have conclusive evidence for Cube in Good that I don't have in Bad—viz., Cube. And since $\text{Prob}(\neg\text{BIV}/\text{Cube}) = 1$, I also have conclusive evidence for not BIV in Good that I don't have in Bad.

We are concerned with the External World instance of EPH—that is to say, with:

- (1) I'm not justified in believing Cube unless I'm justified in believing not BIV.
- (2) I'm not justified in believing not BIV unless something other than the fact that I'm currently having such-and-such sensory experiences as of Cube justifies me in believing not BIV.

- (3) There is nothing else to justify me in believing not BIV.
- (4) ∴ I'm not justified in believing Cube.

It's clear that Williamson would deny the conclusion of this valid argument. But which of its premises would he deny? Given what he says about "intuitive closure" (119), we may infer that Williamson would accept (1). What about (2)? I'm not sure whether Williamson would accept or reject this premise. It presents one problem for him if he accepts it, and another if he rejects it. To see why, we need first to appreciate that, as already intimated, the crux of his reply to the argument will be his denial of premise (3). Here, I believe, he would claim that there is something other than the fact that I'm having such-and-such sensory experiences as of Cube to justify me in believing Cube, because in Good my justification for believing Cube includes evidence which I don't have in Bad—namely, that there is in fact a red cube before me. I'll elaborate on this presently. In the meantime, let's return to the problem (2) presents.

I said that (2) presents one problem for Williamson if he accepts it, and another if he rejects it. The problem in both cases turns on how Williamson can account for my being justified in believing Cube in Bad. (It's clear that it's possible for there to be a token of Bad in which I believe Cube—e.g. one in which I'm only recently envatted—and that when I do I will be justified in believing Cube, unless skepticism is correct and it's impossible for anyone to have a justified perceptual belief about the external world. And, of course, if I'm justified in believing Cube in Bad, then, by the closure principle JCB, I'm also justified in believing not BIV in Bad.) Now, the crux of Williamson's response to the version of skepticism he considers is that one has evidence in good cases that one doesn't have in bad cases precisely by virtue of having perceptual knowledge in good cases. The implication is that this is what is *required* to respond to the skeptic—and presumably to the skeptic about justified belief, as well as to the skeptic about knowledge. If, therefore, Williamson can account for my being justified in believing Cube in Bad in a way that doesn't entail the evidential difference he finds between good and bad cases of perceptual belief, then he will have shown that the response he gives to skepticism in his book isn't *required* to account for how perception can justify us in believing propositions about the external world. That is the problem Williamson would encounter if he denies premise (2), thereby endorsing a claim which entails that no

evidence provided by my *knowing* external world propositions is required for me to be justified in believing not BIV. So it would seem that Williamson is constrained to accept premise (2) on the grounds that I can be justified in believing not BIV only by evidence provided by my knowing external world propositions—in other words, on the grounds that one can't have justified beliefs in bad cases unless one has knowledge in good cases. He has a plausible way of making that case, but only if it's plausible for him to claim that it's impossible for me to believe Cube if I've been a brain in a vat my whole life. We know the twin-earth case to be made for that impossibility; the trouble is that there is some doubt about how good it is when applied to propositions like Cube or BIV. On the other hand, if it's possible only for the recently envatted to believe propositions like Cube, then Williamson might reasonably argue, say, that it's precisely by virtue of knowledge gained in good cases that one can be justified in believing that one is not a brain in a vat.⁹ In any case, I propose that this issue be bracketed for the rest of this paper and that we proceed on the assumption that Williamson can justify his accepting (2) provided he can justify his not accepting (3).

So how, specifically, might Williamson argue against premise (3) of External World? It wouldn't be eristically effective to argue against (3) in a way that presupposed that I was justified in believing Cube, but if Williamson is justified in doubting (3), then he should have at hand an argument to show that the EPH skeptic has not shown herself to be warranted in asserting (3). He would have such an argument if he could argue that the EPH skeptic failed to see that there was something other than the mere fact that I'm having sensory experiences as of Cube that would be available to justify me in believing Cube if I was in Good and thus knew Cube. And it's apt to seem that Williamson has such an argument if his theory of evidence is correct, an argument that in effect showed that his case against *SEA* also provided the wherewithal for a case against *SJA*. The argument to which I allude is as follows:

- (1*) If I'm in Good and thus know Cube, then Cube is
conclusive evidence for me that Cube.

⁹ Nico Silins (2005) exaggerates the extent to which having to account for justified beliefs in bad cases is a problem for Williamson.

- (2*) If Cube is conclusive evidence for me that Cube, then Cube is available to justify me in believing Cube and thus, via JCB, in believing not BIV.
- (3*) \therefore That part of SJA is false which claims that whatever can justify me in believing Cube if I'm in Good and thus really do know Cube will also be available to justify me in believing Cube if I'm in Bad and BIV is true.
- (4*) The EPH skeptic's case for premise (3) of External World presupposes that false part of SJA.
- (5*) \therefore The skeptic is unwarranted in asserting premise (3) (since her case for it relies on an unwarranted false assumption).

Let's call this argument W. Might the EPH skeptic have a way to question W?

I think she may; I think she might well have doubts about W's premise (2*). Given that E is conclusive evidence for P for me if my evidence includes E and $\text{Prob}(P/E) = 1$, then it does follow from Williamson's theory that Cube is conclusive evidence for Cube for me. And it does seem right that if E is conclusive evidence for P for me, then E is a conclusive justification that I have for believing P. But in distinguishing the three justification relations IS, HAS, and IN, we saw that E can be a justification that one has for believing P but yet not be what justifies one *in* believing P, even when one is justified in believing P. So, while Williamson's theory of evidence might entitle him to claim that Cube is a conclusive justification that I *have* for believing Cube, given that I know Cube, he hasn't *thereby* shown that Cube is available to justify me *in* believing P, given that I know Cube. Even when we grant his theory of evidence, we may still question whether Williamson is in a position to claim that Cube is something that may justify me *in* believing Cube, given that I know Cube.

But isn't it *analytic* that if E is conclusive evidence for P for me, then E at least stands *available*, as things are, to justify me *in* believing P? Actually, it isn't true, let alone analytic, that E stands available to justify me in believing P if E is conclusive evidence for P for me. For suppose I see that my patio is wet and on the basis of that

evidence come to know, and thus to be justified in believing, that my patio is wet because it rained during the night. If Williamson's EV and $E = K$ are correct, then the fact that my patio is wet because it rained during the night is *conclusive evidence* that my patio is wet; in fact, evidence doesn't get any better than that. But the fact that my patio is wet because it rained during the night isn't—and can't be, given the facts of the story—what *justifies me in believing* that my patio is wet. What justifies me in believing that my patio is wet, and all that is available in the circumstances to justify me in believing that my patio is wet, is that I saw that it was.

In fact, Williamson himself implicitly acknowledges that E's being conclusive evidence for x for P doesn't secure that E stands available to justify x in believing P. The following passage leaves little doubt that Williamson would agree that the awkward symmetry noted in the patio example presents a counterexample to the claim that if E is conclusive evidence for P for me, then E is available to justify me in believing P:

If all knowledge is evidence, then EV ... [has] the effect of making evidential interconnections within one's knowledge symmetric. For $\text{Prob}(P/Q) > \text{Prob}(P)$ if and only if $\text{Prob}(P \& Q) > \text{Prob}(P)\text{Prob}(Q)$; since the latter condition is symmetric in P and Q, $\text{Prob}(P/Q) > \text{Prob}(P)$ if and only if $\text{Prob}(Q/P) > \text{Prob}(Q)$. Thus, given that S's evidence includes both P and Q, P is evidence for Q for S if and only if Q is evidence for P for S by EV. Consequently, given that one knows P and Q and that all knowledge is evidence, EV implies that if P is evidence for Q for one then Q is evidence for P for one. We could avoid this result by modifying EV. For example, we could stipulate that E is evidence for H for S only if S's belief in E does not essentially depend on inference from H. But it might be neater to retain EV unmodified and say that E is *independent* evidence for H for S only if S's belief in E does not essentially depend on inference from H. (204)

If we accept Williamson's theory of evidence, then in Good Cube is conclusive evidence for Cube for me. But we've just seen that even Williamson must admit that something can be conclusive evidence for P for me yet incapable of justifying me in believing P. So, given that we accept Williamson's theory of evidence and given that E is conclusive evidence for P for me, *what else* must be true of E in order for it to be able to justify me in believing P?

Williamson's "independent evidence" fix suggests that he would say:

E justifies S in believing P *only if* E is *independent*
evidence for P for S,

from which it follows that

Even if E is extremely strong or conclusive evidence for P
for S, E is incapable of justifying S in believing P if E is
not independent evidence for P for S.

(E is *independent evidence* for P for S iff (i) E is evidence for P for S and (ii) S doesn't believe E on the basis of P ("S's belief in E does not essentially depend on inference from" P).)

Can we also say that

E is capable of justifying S in believing P *if* E is extremely
strong or conclusive *independent* evidence for P for S,

or are there still further necessary conditions that extremely strong or conclusive independent evidence must satisfy if it's to be capable of justifying S in believing P? I believe that the following three examples show that the displayed condition is false and that, therefore, some further condition is required.

Raven example. Suppose I know that the next observed raven will be black on the basis of knowing that all observed ravens have been black. It's surely preposterous to say that even part of what justifies me in believing that the next observed raven will be black is that the next observed raven will be black (what justifies me in believing that the next observed raven will be black is that all observed ravens have been black). But the fact that the next observed raven will be black is for me conclusive independent evidence that the next observed raven will be black.

Smithers example. I know (because his instructor told me) that Smithers failed his logic final and on that basis know, and am justified in believing, that a D is the best grade he can receive in the course [D, for short]. I also independently know that Smithers didn't study for the final, and that is pretty good evidence for me that D. But there's this asymmetry between the two evidence facts. The fact that Smithers failed the final has its evidential status for me regardless of whether or not Smithers studied for the final: I can infer D from that whether or not I even believe that he didn't study, but I couldn't infer D from the fact that he didn't study unless my reason for believing that he failed the final was just that he didn't study for it. Here the fact that Smithers didn't study may be strong enough independent evidence for D for me—independent because I didn't infer that he didn't study from D—but, nevertheless, incapable in the circumstances of justifying me in believing D because of the way it's screened off from the only thing in the circumstances that could justify me in believing D—viz., the fact that Smithers failed the final.

Coke example. I infer, and thereby come to know, that the Coke machine is sold out from the fact that the machine's "Sold Out" sign is lit. I would be justified in inferring that the machine is sold out from the fact that it says it's sold out whether or not the machine is sold out; but, as I've no other way in the circumstances of inferring that the machine is sold out, I justifiably wouldn't believe that the machine was sold out unless I inferred that from the fact that the machine says it's sold out. Now, in the circumstances, the fact that the Coke machine is sold out is conclusive independent evidence for me that the Coke machine is sold out, but that isn't what justifies me in believing that the machine is sold out. What justifies me in believing that is that the machine says it's sold out. So, once again, we see that E can be conclusive independent evidence for P for S yet incapable of justifying S in believing P.

Let me suggest then the following criterion, which assumes that Williamson's theory of evidence is correct (and which uses 'infer' in Williamson's sense, a sense some might think is better expressed by 'on the basis of', especially as regards the way in which sensory experiences function to justify the beliefs they induce¹⁰):

¹⁰ Many philosophers think that, e.g., a creature's having a visual experience as of P can justify it in believing P even though it doesn't know, or even believe, that it's having that

E ≠ J: Even if E is very strong or conclusive independent evidence for P for S in circumstances C, E is incapable of justifying S in believing P in C if in C there is evidence E' such that (i) S can become justified in believing P in C only by inferring P from E', and (ii) S's becoming justified in believing P in C by inferring P from E' doesn't depend on E's being true.

Thus, in the raven example, E = P = the fact that the next observed raven will be black, and E' = the fact that all observed ravens have been black; in the Smithers example, E = the fact that Smithers didn't study for the final, E' = the fact that Smithers failed the final; and P = the proposition that a D is the best grade Smithers can receive in the course; and in the Coke example, E = P = the proposition that the Coke machine is sold out, and E' = the fact that the machine says that it's sold out.

Now, by definition of Good, I know Cube (= that there is a red cube before me) in Good and don't infer Cube from Cube, and therefore, given Williamson's theory of evidence, Cube is conclusive independent evidence for Cube for me in Good. But if E ≠ J is correct, Cube isn't available in Good to justify me in believing Cube. For (i) in Good I come to be justified in believing Cube by inferring it from the fact that I'm having such-and-such sensory experiences as of Cube, and I can't become justified in believing Cube in Good other than by inferring it from that evidence; and (ii) my becoming justified in believing Cube by inferring it from the fact that I'm having those sensory experiences doesn't depend on Cube's being true—I would become justified in believing Cube by that inference even if all else were the same except that Cube was false.

So I provisionally conclude that, while Williamson may have given us reason to disbelieve *SEA*—the assumption that I have the same *evidence* for Cube in Good and in Bad—he hasn't given us good reason to disbelieve *SJA*—the assumption that what *justifies me in believing* Cube in Good is the same as what justifies me in believing Cube

visual experience. Williamson must deny this, since he holds that E = K and that only evidence can justify a belief. He argues against the claim that perceptual experience is a kind of non-propositional evidence on pp. 197-200.

in Bad. Applied to the argument W—the argument one might think one discerns in Williamson—we see that Williamson hasn't entitled us to think that W is sound, because, even if we concede his theory of evidence and allow that in Good Cube is conclusive evidence for Cube for me, he still hasn't entitled us believe that premise (2*) is true. Further, since it is plausible that one knows a proposition only if one is justified in believing it, Williamson also hasn't shown that it's even possible for me to be in Good.

IV. Some Possible Replies

I reckon the probability that Tim Williamson will accept my argument to show that he hasn't provided a solution to EPH skepticism to be, say, ≤ 0.000013 . But *how* will he respond to it? If I've correctly represented how he would respond to EPH skepticism, he must either deny that the application of $E \neq J$ to Good shows that Cube is incapable of justifying me in believing Cube in Good, or else he must deny the criterion $E \neq J$.

There are in principle two ways to deny my claim about the application of $E \neq J$ to Good. One might deny that condition (i) is satisfied by arguing that I can become justified in believing Cube in Good in some way other than by inferring it from Experience (= the fact that I'm having such-and-such sensory experiences as of Cube), or one might deny that condition (ii) is satisfied by arguing that Cube does have to be true in order for me to become justified in believing Cube in Good by inferring it from Experience. Both ways seem unpromising to me.

The first way requires my becoming justified in believing Cube in some way other than by inference from Experience. What could such a way possibly be? It can't be that one becomes justified in believing Cube in Good by *inferring* it from Cube. If one did infer Cube from itself, then Cube would not be *independent* evidence for itself, and thus, evidently, ruled out on that account as being that which justifies me in believing it. There are cases where it's perhaps not unreasonable to suppose that the fact that P justifies believing P. One might hold that what justifies one in believing that one is in pain is just the fact that one is in pain. But, in the first place, perceptual beliefs don't seem at all like that, and, in the second place, Williamson seems not to be in a position felicitously to hold even that the fact that Sally is in pain justifies her in believing that she is in pain. For what is not altogether implausible is that what justifies Sally in believing

that she is in pain is just the fact that she is in pain; but it doesn't seem at all plausible that what justifies Sally in believing that she is in pain is that she knows that she is in pain and infers that she is in pain from the fact that she is in pain. One can hardly become justified in believing a proposition by inferring it from itself. I suppose Williamson would have to say that Sally has underived knowledge that she is in pain and that that is what makes her justified in believing that she is in pain. To reconcile this with his doctrine that only evidence can justify, he could say that the fact that she is in pain justifies her in believing that she is in pain somehow by virtue of its being conclusive evidence that she is in pain, but not by virtue of her *inferring* that she is in pain from that evidence. One wants to hear more, but in any case the model doesn't fit my knowing Cube by perception.

The second way to deny my claim about the application of $E \neq J$ to Good—denying that condition (ii) is satisfied—strikes me as even more unpromising. If I'm not justified in believing Cube in Bad, it will be because the EPH skeptical argument is sound. If it's possible for me to be justified in believing Cube at all, then I'm surely justified in believing it in Bad. Should I learn that BIV was true, I certainly wouldn't conclude that I wasn't justified in believing Cube. But if I'm justified in believing Cube in Bad, then it's surely by inference (given Williamson's theory of evidence, which I'm taking as given). So, as Cube is false in Bad, Cube's being true can't be a necessary condition of my becoming justified in believing Cube by inferring it from Experience. And if it's not a necessary condition in that way in Bad, then it's very implausible that something about Good makes it a necessary condition in Good.

So much for denying the application of $E \neq J$ to Good. Perhaps denying the criterion $E \neq J$ will yield a more promising response. There are, after all, *prima facie* counterexamples to $E \neq J$. For example, when asked what justifies him in thinking that Alice kissed Ben, it might be appropriate for Harold to reply that Alice informed him that she kissed Ben. Harold would not have come to believe that Alice kissed Ben by inferring that Alice kissed Ben from the proposition that Alice informed him that she kissed Ben, since in order for him to believe that Alice *informed* him that she kissed Ben he would already have to believe that Alice kissed Ben. This is apt to appear to be a counterexample to $E \neq J$ because the fact that Alice informed Harold that she kissed Ben

is conclusive evidence that she kissed Ben (x informed y that P entails P), but in the circumstances Harold could become justified in believing that Alice kissed Ben only by inferring that she did from the evidence that she *told* him that she kissed Ben, and Harold's becoming justified in that way in believing that Alice kissed Ben doesn't depend on its being true that Alice kissed Ben.¹¹

It's unclear whether the example provides a counterexample because being asked what justifies someone in believing a proposition is a request for an explanation, and we often appropriately respond to such questions in ways that don't actually give the correct explanation or give the explanation embedded in information that isn't essential to the explanation, as when we explain that the car won't start because something's wrong with the ignition, or that the window broke because your niece Wilma kicked her new orange soccer ball into the window. Still, a more systematic way of challenging $E \neq J$ might proceed in the following way:

We need to distinguish *the way in which* x becomes justified in believing P from *that which* justifies x in believing P. The idea is that while I become justified in believing Cube in the same way both in Good and in Bad—viz., by inferring Good from Experience—the justification I acquire in that way in Good differs crucially from that which I acquire in Bad. Roughly speaking, the justification I acquire in Good includes the justification I acquire in Bad, but has as an additional component the fact that Cube. I gain one justification for believing Cube when I infer it from Experience, but I gain an even better one when the fact that I had such-and-such sensory experiences as of Cube was caused by the fact that there was a red cube before me. Both in Good and in Bad I become justified in believing Cube by inferring it from Experience, but in Bad the justification I have for believing Cube consists just in the fact that I had such-and-such sensory experiences as of

¹¹ This sort of example was pressed on me by Anna-Sara Malmgren and Nico Silins.

a red cube, whereas in Good it also contains the additional evidence for Cube that is owed to the fact that my having such-and-such sensory experiences as of a red cube was caused by the fact that there was a red cube before me.

There may be more than one thing wrong with this response, but the main thing wrong with it is that (a) I wouldn't be justified in being more confident of Cube in Good than I am in Bad, but (b) I would be so justified if I had a better justification, one based on better evidence, for believing Cube in Good than I have for believing it in Bad.¹² Let me motivate (a) by starting with a change of example.

I wake up one February morning and look out the window. I see that Washington Square Park and the surrounding streets are covered in snow, and I infer that it snowed during the night. I have no positive reason, skepticism aside, to suspect that it didn't snow, but of course I'm aware of the *possibility* of various ways in which, compatible with what I seem to see, it didn't snow during the night, and thus the degree of confidence that I'm justified in having is less than complete confidence, though still pretty high. Let's pretend that degrees of confidence can be measured by real numbers in the interval $[0, 1]$ and suppose I'm justified in being confident to degree 0.93 that it snowed during the night. Now, my description of the scenario is compatible with two more complete descriptions of it. In one, it did snow during the night, and that is why I see the snow. In this completion, I count as knowing that it snowed during the night. In the other completion, though I had no reason at all to suspect it, the snow didn't fall from the sky but was artificially manufactured and placed on the ground by a film crew which wasn't visible when I saw the snow. By construction of the example, I remain confident to degree 0.93 that it snowed in both completions. I submit that my being justified in having that degree, but no greater degree, of confidence is unaffected by which completion obtains. The parallel with the ongoing Cube example should be obvious. I seem to see a red cube before me and I have no reason (skepticism aside) to doubt that my experience is veridical other than my knowledge of the possible ways in which it might not be. So I'm not justified in being absolutely confident that there's a red cube

¹² This objection derives from the "equal justification" problem Nico Silins (2005) raises for Williamson's theory of evidence.

before me, but I'm justified in being pretty confident. Let's say I'm confident to degree 0.93. I submit that I remain justified in having that degree, but no greater degree, of confidence in Cube whether the completing description of my situation places me in Good or in Bad.

So much for (a). Might (b) be denied? If my justification for believing Cube in Good is better than my justification for believing it in Bad, then I am more justified in believing Cube in Good than I am in Bad. By construction of the example, my degree of confidence in Cube is the same in both Good and Bad. If I'm more justified in believing Cube in Good than I am in Bad, then I *should* be more confident of Cube in Good. If I'm less confident than I should be, then I'm not justified in having the degree of confidence I have. As Silins makes clear,¹³ it would be rather bizarre to hold that my degree of confidence in Cube is justified in Bad but not in Good. I see no reason to suppose I'm doing anything epistemically wrong in Good. On the contrary, what would be epistemically wrong would be for me to be more confident of Cube in Good than I am.

I suppose Williamson must disagree. After all, he says that, for any proposition P, if you know P, then P has evidential probability 1 for you, and he also says that "rationality requires one to conform one's beliefs to one's evidence" (12), where by this he means that "the norm of credence is to proportion one's degree of belief to the evidence."¹⁴ This, I take it, means that rationality requires me to believe Cube to degree 1 in Good. Still, Williamson is explicit about our needing "a conception of rationality on which we are not always in a position to know what it demands" (15). So perhaps he would say that, while I'm being perfectly reasonable in believing Cube only to degree 0.93 in Good, this is because I am not in a position to know what rationality requires of me. But it's not clear that he can say that. When Williamson acknowledges that we aren't always in a position to know what rationality requires of us, in requiring us to respect our evidence, it's because we aren't always in a position to know what our evidence is. Now, for Williamson, who holds that $E = K$, if I don't know what my evidence is, this can only be because I don't know what I know. But in a normal case of Good, I would not only know Cube, but would also know that I know Cube, and

¹³ Silins (2005).

¹⁴ Williamson (2005), p. 432.

therefore I *would* know what my evidence is, and thus, presumably, *would* be in position to know what rationality requires of me. Could he argue that I might know what my evidence is but still not know what rationality requires of me, either because I don't know that the probability of Cube on my evidence is 1 or that I don't know that "the norm of credence is to proportion one's degree of belief to the evidence"? I doubt it. I doubt that Williamson would want to say that the only people who are in a position to know what rationality requires of them are those who accept his theory of evidence and rationality. In believing Cube to degree 0.93 in Good, I seem neither to be acting irrationally nor failing to know what rationality requires of me. I see no reason to doubt that appearance.¹⁵

¹⁵ Earlier versions of this paper were given as talks at Rutgers University and the University of St. Andrews, and presented in a seminar I gave at NYU, all in fall 2005. The final version benefited from the discussions at those events, and from written communications or comments from Adam Elga, John Hawthorne, Anna-Sara Malmgren, Nico Silins, and Dean Zimmerman.

References

- Burge, T. (2003). "Perceptual Entitlement," *Philosophy and Phenomenological Research* 67: 503-548
- Dretske, F. (1970). "Epistemic Operators," *Journal of Philosophy* 67: 1007-23.
- Martin, M. (2004). "The Limits of Self-Awareness," *Philosophical Studies* 120: 37-89.
- McDowell, J. (1998). "Criteria, Defeasibility, and Knowledge," in *Meaning, Knowledge, and Reality* (Harvard University Press).
- Nozick, R. (1981). *Philosophical Explanations* (Harvard University Press).
- Peacocke, C. (2004). *The Realm of Reason* (Oxford University Press).
- Pollock, J. (1986). *Contemporary Theories of Knowledge* (Rowman and Littlefield).
- Pryor, J. (2000). "The Skeptic and the Dogmatist," *Noûs* 34: 517-49.
- Silins, N. (2005). *Philosophical Perspectives* 19: 375-404.
- Vogel, J. (1990). "Cartesian Scepticism and Inference to the Best Explanation," *Journal of Philosophy* 87: 658-64
- White, R. (forthcoming). "Problems for Dogmatism," *Philosophical Studies*.
- Williamson, T. (2000). *Knowledge and Its Limits* (Oxford University Press).
- _____ (2005). "Précis of *Knowledge and Its Limits*," *Philosophy and Phenomenological Research* 70: 431-5.
- Wright, C. (2004). "Warrant for Nothing (And Foundations for Free)?" *Supplement to the Proceedings of the Aristotelian Society* 78 (2004): 167-75.