

Evidential Externalism

Jeffrey Dunn

1 Introduction

There are at least two interesting questions about evidence that one might ask. First, one might ask:

When, and under what conditions, is proposition P evidence for proposition Q ?

We could call this the *Evidence-For Question*. A different question one might ask is:

When, and under what conditions, is proposition P evidence (for agent S)?

This is a different question. Call this the *Evidence Question*. The Evidence-For Question has received significant attention. Debates in the Bayesian literature about the correct confirmation function, for instance, are aimed at the answering the Evidence-For Question.¹ The Evidence Question is somewhat less investigated.² Nevertheless, it is important. According to prominent epistemological views, an agent's beliefs should be responsive to his evidence. Bayesian epistemology presents one precise account of how this goes: agents are to update their beliefs by *conditionalizing* on the evidence received. Without an answer to the Evidence Question, dictums like this are empty.

Nicholas Silins ([2005]) has recently offered a partial answer to the Evidence Question. In particular, Silins argues for *Evidential Internalism* and against *Evidential Externalism*. According to Silins, Evidential Internalism is the following thesis:

EI: Necessarily, if A and B are internal twins, then A and B have the same evidence. (p. 376)

¹See, for instance, Fitelson ([2001]).

²Though, see Maher ([1996]), Williamson ([2000]), and Neta ([2008]).

Silins maintains that A and B are internal twins just in case they are the same in all their non-factive mental states.³ So, Silins holds that agents with the same non-factive mental states are in possession of the same evidence.

Silins argues for EI by arguing against its negation, Evidential Externalism:

EE: Possibly, internal twins A and B have different evidence.

An answer to the Evidence Question according to which an agent's evidence depends on the reliability of various belief-forming processes, is an example of an account of evidence committed to EE. Another answer to the Evidence Question committed to EE is Williamson's ([2000]) $E = K$ thesis, which states that an agent's evidence consists of all and only the propositions that the agent knows. These answers have attractive features, so an argument ruling out *any* answer committed to EE is of significant interest. Indeed, the argument is of even more interest than this would suggest. Debates about the right answer to the Evidence Question are really debates about the appropriate evidential ground for belief. If Evidential Externalism is ruled out, we go a long way towards saying what the appropriate evidential grounds are.

Silins's main argument against EE turns on showing that it is a consequence of EE that radically deceived agents are sometimes more justified than non-deceived agents in very similar circumstances. This is thought to be an intolerable consequence. In this paper I will defend EE by responding to this argument. The response reveals something interesting about how to answer the Evidence Question and indeed about evidence itself.

Here, then, is where we're headed. I will first explain Silins's argument. After this I will give two responses. First, I will argue that the alleged unacceptable consequence of EE is not always unacceptable. Second, I will show that EI is subject to an argument structurally identical to the one Silins gives against EE. These two responses work together: the less convinced one is by the first response, the more pressure there is to accept the second response. Finally, after giving these responses, I will explain what can be learned about how to go about answering the Evidence Question.

Before beginning, I should note that I will consider Silins's argument within the framework of standard Bayesian epistemology. So, I'll help myself to the following: First, I'll only be considering evidence that is propositional.

³"On the use of 'internal' I follow, two people are internal twins just in case they have the same non factive mental states to the same degree—the same beliefs, apparent experiences, apparent memories, and so on." (pp. 376-7).

Second, I'll assume that beliefs come in degrees so that a total doxastic state can be represented by a credence function, $cr(\cdot)$, and that rational credence functions are probabilistically coherent.⁴ Third, to ease exposition, I'll use the term 'evidence set' to refer to the set of propositions that are evidence for an agent. Finally, I'll assume that Conditionalization is the principle that dictates how evidence should be accommodated. Where E is the conjunction of propositions in your evidence set, ' cr_E ' is your credence function with that evidence, and ' cr ' your credence function without that evidence, Conditionalization says that $cr_E(P) = cr(P|E)$. An important consequence of this is that $cr_E(E) = cr(E|E) = 1$. Thus, a subject's evidence receives credence 1.

None of these assumptions lead to a misrepresentation of Silins's argument. Silins, too, only countenances evidence that is propositional, and assumes that beliefs come in degrees that can be represented by degrees of confidence. Though he doesn't explicitly endorse Conditionalization, all of his examples are such that a subject's evidence should receive full credence. Assuming Conditionalization is just one way of making it clear why this is so.⁵

2 Silins's Argument

Silins's argument attempts to show that EE has the implausible consequence that a radically deceived agent can sometimes be epistemically better off than his non-deceived counterpart.

Here is the scenario meant to establish that EE has this consequence. Gary is in the good case. He seems to remember having a banana yesterday, and he is right about this. Let $B =$ "I had a banana for breakfast yesterday." Imagine that B is true when considered by Gary. Further, assume that EE is true and that the right external relation obtains between Gary and the fact that B so that B is a member of Gary's evidence set. Perhaps

⁴By this I mean that rational credence functions obey the standard Kolmogorov axioms:

1. For all P , $cr(P) \geq 0$,
2. For all logical truths, \top , $cr(\top) = 1$, and
3. $cr(P \vee Q) = cr(P) + cr(Q)$ for all P, Q such that $(P \wedge Q)$ is contradictory.

I also understand conditional degrees of belief in the usual way: $cr(P|Q) =_{df} cr(P \wedge Q)/cr(Q)$.

⁵For those concerned about the requirement that evidence receive credence 1, see the Appendix where I show that the arguments—both Silins's and my own—do not depend on this assumption.

Gary knows that B . Or perhaps Gary remembers that B , and memory is sufficiently reliable to render B a member of his evidence set. Let SB = “I seem to remember that I had a banana yesterday.” We assume that SB is true when considered by Gary, and that the right relation obtains between Gary and the fact that SB so that SB is a member of Gary’s evidence set. So, both B and SB are members of Gary’s evidence set according to EE. However, realizing his memory can sometimes be mistaken Gary is such that $cr_E(B) = 0.9$, while $cr_E(SB) = 1$.

Now consider Barry, an internal twin of Gary, who is unfortunately nothing more than an envatted brain. Barry is in a bad case: he didn’t actually have a banana, although he seems to remember doing so. Since Barry didn’t actually have a banana yesterday, B is false when considered by Barry. Accordingly, the right external relation does not obtain between Barry and the fact that B so B is not a member of Barry’s evidence set. Nevertheless, as he is an internal twin of Gary, Barry does seem to remember having a banana. Since all the internal facts about Gary are true of Barry, SB is a member of Barry’s evidence set. Further, Barry is such that $cr_E(B) = 0.9$ and $cr_E(SB) = 1$, again because they are internal twins.⁶

To fill in the scenario we can stipulate that that Barry and Gary are both such that $cr(B|SB) = 0.9$. Then we get the following result: Barry has the credences that result from conditionalizing on Barry’s evidence, but Gary doesn’t have the credences that result from conditionalizing on Gary’s evidence. Gary’s $cr_E(B) = 0.9$, even though B is evidence for Gary. Thus, Barry is doing epistemically better than Gary in this scenario, even though Barry is in the bad case, and his internal twin Gary is in the good case.⁷

We can summarize Silins’s argument as follows:

⁶Note that B and SB are both propositions which feature the indexical expression ‘I’. By characterizing Barry and Gary as having beliefs in propositions like B and SB , Barry and Gary seem like internal twins. Of course, there is a sense in which when Barry believes B he is believing something very different than when Gary believes B . I note this complication only to set it aside, as it doesn’t seem to play any role in the discussion. For our purposes, if Gary and Barry both believe B and SB to the same degree (and everything else about their doxastic states are the same) then they have identical doxastic states.

⁷If we grant an assumption Silins makes about degrees of justification, we can show that Gary is less justified than Barry. Silins writes: “. . . if one’s degree of confidence ought to be n , and one’s actual degree of confidence diverges from n , then one’s actual degree of confidence is less than fully justified insofar as it diverges from what it ought to be.” (p. 387) Since Gary’s credence in B ought to be 1 (since it is evidence, and $cr(B|B) = 1$), Gary is not fully justified. Since Barry’s credence in B ought to be 0.9 (since SB is Barry’s only evidence, and $cr(B|SB) = 0.9$), Barry is fully justified. So, Gary is less justified than Barry.

1. If EE is true, then a subject in the bad case is sometimes epistemically better off than his internal twin in the good case.
 2. A subject in the bad case is never epistemically better off than his internal twin in the good case.
-
- C. Thus, EE is false.

Premise 1 is meant to be established by the scenario of Gary and Barry. Gary is not deceived in any way, his internal twin Barry is radically deceived, so Barry is in the bad case and Gary is in the good case. Nevertheless, as we just saw, EE seems to say that Barry is doing epistemically better than Gary.

A defender of EE could challenge this. After all, the defender of EE need not say that *every* external difference makes an evidential difference. Perhaps there are features of *this* particular case the defender of EE could use to show that Gary's evidence does not include *B*. However, Silins claims that if EE is to be an interesting thesis then there are bound to be *some* cases where the agent in the good case has *more* evidence than the agent in the bad case. If the agent in the good case doesn't respond to this evidence, then we have a situation where the bad case is better. This would establish premise 1.

It is worth noting that things are probably not as simple as this. There are, after all, possible EE views according to which agents in the good case will not have *more* evidence than their counterparts in the bad case, but instead will have completely *different* evidence. For example, Matthew Kennedy ([2010]) has recently considered Silins's argument from the perspective of Naive Realism about perception. Since Naive Realism is a thesis about perception, the relevant Silins-style argument is not one that concerns beliefs based on *memory*, but rather beliefs based on perceptual experiences. So, consider a case where Gary sees a banana before him, but Barry only hallucinates the banana. The kind of naive realism that Kennedy defends says that Gary has different evidence than Barry in virtue of differences in their experiences. Gary has something like: "I see the banana" as evidence, whereas Barry has something like *H*: "I am hallucinating as of a banana" as evidence. This is to opt for a version of EE according to which the agent in the good case doesn't have *more* evidence, rather, he has *different* evidence. Arguably, such a version of EE could be interesting, and would allow one to deny premise 1. This, in fact, is Kennedy's response to the argument.

Let me note how my response will differ from Kennedy's. The response that I will give is not committed to naive realism, nor the specific claims Kennedy makes about the differences in evidence that naive realism gives

rise to. Many defenders of EE (for instance, Williamson [2000]) do not rest their commitment to EE on naive realism, and they accept that sometimes the evidence set in the bad case *is* a proper subset of the evidence set in the good case. Kennedy’s response is not applicable to such defenders of EE, whereas my response is.⁸

At any rate, I will take premise 1 for granted. The reason for doing so is that most, though not all, versions of EE are such that premise 1 is true. It is true because most versions of EE imply that the agent in the good case can have *more* evidence than the agent in the bad case.

The support for premise 2 is more sparse. Silins claims that premise 2 is an instance of a very plausible principle:

The Bad Case is Never Better (BCNB): Necessarily, if B is in the bad case and A is an internal twin of B in the good case, B is not more justified in believing P than A. (p. 389)

One might take the story about Gary and Barry to provide some support for BCNB. Perhaps it is supposed to be obvious from the story that Gary is not epistemically worse off than Barry, and so provide some inductive support for the principle. In the next section, I will present a scenario designed to cast doubt on BCNB. If BCNB is mistaken, then premise 2 is unsupported and Silins’s argument fails.

Before doing this, however, one might wonder why the argument requires such a strong premise like BCNB.⁹ After all, it seems that there might be a more modest argument against EE, of the form:

- 1*. If EE is true, then Barry (in the bad case) is epistemically better off than Gary (in the good case).
 - 2*. Barry is not epistemically better off than Gary.
-
- C. Thus, EE is false.

This argument need not appeal to BCNB. So, the objection goes, one must do more to respond to Silins’s argument than simply show BCNB false.

There are, however, at least three problems with this modified argument. First, it is clearly not Silins’s argument. Silins is clear that the case of Barry

⁸I also note that for Kennedy’s response to work, his understanding of how evidence affects beliefs must be non-standard. This is because Kennedy is explicit that he does not want the evidence in the bad case to allow the agent to *know* that he is hallucinating. But if H is the evidence in the bad case, then $cr_E(H) = 1$, since $cr(H|H) = 1$, and the agent in the bad case *can* know that he is hallucinating. Thus, Kennedy’s response to Silins is not one amenable to the traditional Bayesian picture about evidence. My response does not have this feature.

⁹Thanks to an anonymous referee for emphasizing this point.

and Gary violates the strong principle, BCNB. Indeed Silins emphasizes this when he writes:

Even if one is willing to accept the result that one is sometimes more justified in the good case, it is harder to live with the claim that one is sometimes more justified in the bad case. [...] I take it that no ordinary reliabilist will be happy to say that one is sometimes better off in the bad case. The consequence that one is sometimes more justified in the bad case should be unwelcome to all. (p. 390)

Second, if the modified argument is to be successful, then we need some reason to believe premise 2*. One could simply say that it is intuitively clear that Barry is not epistemically better off than Gary in the case presented. But, why not? Where does this intuition come from? Clearly, what drives this intuition is something like BCNB: radically deceived internal twins are *never* epistemically better off than their non-deceived counterparts. But if it is BCNB that supports premise 2* of the revised argument, then showing BCNB false *is* relevant, even to this revised argument.

The third problem with this revised argument concerns premise 1*. For, as noted above, there may be features of this specific case that the defender of EE can use to explain why Barry is not epistemically better off than Gary. Perhaps Gary's doubts about his memory are sufficient to make *B* not part of Gary's evidence. The defender of EE could adopt such an account since EE does not imply that every external difference results in an evidential difference. This would show premise 1* false. In the original argument, it is different. Even if the case of Gary and Barry is objectionable to the defender of EE for some reason, still it must be admitted that there are *some* scenarios where the defender of EE will have to say the agent in the bad case is doing better. This is only objectionable, however, if something like BCNB is true.

Thus, the argument against EE really is the original argument. In the next section, I will construct an example designed to show BCNB false.

3 Against BCNB

BCNB says that an agent in the bad case is never doing better than his internal twin in the good case. As mentioned above, this principle gains plausibility by considering the case of Gary and Barry. In that case it seems as though Gary, quite reasonably, has modest doubts about his memory.

Thus, Gary is doing nothing wrong in according less credence to B than to SB . But if Gary is doing nothing wrong, then it is impossible that brain-in-a-vat Barry could be doing *better*. In what follows, however, I will show that there are cases where the doubts that Gary has do not look so appropriate. Such a case will show that sometimes an agent in the bad case *is* doing better than his internal twin in the good case.

Here is one such case. Let R = “there is a red stop sign in front of me.” Assume that there is a red stop sign in front of Gina. She’s looking right at it, and the lighting is ideal. Gina has never been wrong about the presence or absence of stop signs in such conditions, and she doesn’t have any reason to entertain skeptical doubts. Let SR = “I seem to see a red stop sign.” SR is also true of Gina. Stipulate that the relation that Gina bears to R is the same as the relation that she bears to SR , with the exception that SR is something she shares with her internal twins, whereas R is not. (For instance, we could suppose that Gina *knows* both R and SR). I maintain that it is plausible in this situation that both R and SR are members of Gina’s evidence set. Certainly, a defender of EE will find this plausible. However, we stipulate that in this situation Gina is such that $cr_E(R) = 0.9$ and $cr_E(SR) = 1$.

Beth, an internal twin of Gina, is an envatted brain. Given this, there is no red stop sign in front of Beth. Nevertheless, as she is an internal duplicate of Gina, Beth does seem to see a stop sign. Now, Beth does bear the same relation to SR as Gina does. However, in virtue of her deficient situation, Beth does not bear the same relation to R as Gina does. Beth is not in an ideal situation to detect R , the lighting is certainly not ideal, and Beth is repeatedly wrong about propositions like R . Accordingly, while SR is a member of Beth’s evidence set, R is not. Since Gina and Beth are internal twins, Beth is such that $cr_E(R) = 0.9$ and $cr_E(SR) = 1$.

Now, let’s stipulate that that Beth and Gina are both such that $cr(R|SR) = 0.9$. But then we get the following result: Beth has the credences that result from conditionalizing on her evidence, but Gina doesn’t. Gina’s $cr_E(R) = 0.9$, even though R is evidence for Gina. Thus, Beth is doing better than Gina in this scenario, even though Beth is in the bad case, and Gina is in the good case.

I submit that in this case Gina really is doing epistemically worse than Beth.¹⁰ The reason for this is that this is not a situation in which Gina

¹⁰I maintain that Gina is doing epistemically worse *in the sense that matters to BCNB*. There is a sense in which Gina is doing better in that her beliefs are closer to the truth. Gina’s credence in R is 0.9 and R is true, whereas Beth’s credence in R is 0.9 and R is false. However, this can’t be what is intended by BCNB, since this would imply that Gary

should be doubtful with respect to R . Further, there's a clear sense in which Gina is doing worse here. There is a red stop sign directly in front of Gina, the environment is ideal, and yet Gina doesn't treat R as evidence. There is a clear sense in which Gina would be doing better if she believed R to degree 1—it follows from conditionalizing on her evidence. Further, in this same sense, Beth is fully justified in believing R to degree 0.9. The environment is not ideal for Beth, she has only SR as evidence, and so conditionalizing on this results in $cr_E(R) = 0.9$. Thus, if Gina were such that $cr_E(R) = 1$, this would be just as justified as Beth's $cr_E(R) = 0.9$. Now, since Gina and Beth are internal twins, they cannot give different values to $cr_E(R)$. Altering Gina's $cr_E(R)$ from 1 to 0.9, makes Gina and Beth internal twins. But surely altering Gina's belief state in this way doesn't change Beth's degree of justification. So, in doing this, Gina's degree of justification must change. But it can't change for the better, for she was fully justified to begin with, so it must change for the worse. So, in the scenario, it is appropriate to say that Gina is less justified than her brain-in-a-vat twin, Beth. BCNB says this can't happen, so it is false.

I anticipate objections to my claim that this example shows BCNB to be false. I'll address these momentarily. Before that, however, note three things.

First, it is natural to have worries about this example that have nothing to do with the issues separating EI from EE, but rather having to do with the framework according to which evidence receives credence 1. One might think that *no* propositions should receive full credence, let alone a proposition like R . This issue, however, is tangential to the issue dividing EE and EI. In the Appendix I take up this issue in detail for those who are interested.

Second, in Section 4 of this paper I make a different response to Silins's argument. This response assumes that one is *not* convinced of the falsity of BCNB by the case of Gina and Beth. So if one finds oneself unconvinced by this example, hold on until that part of the paper.

Finally, note that any attempt to defend BCNB against the Gina and Beth counterexample cannot simply appeal to what Silins calls *Equal Justification*:

Equal Justification: Necessarily, if A and B are internal twins, then A is justified to degree n_1 in believing P to degree n_2 just in case B is justified to degree n_1 in believing P to degree n_2 . (p. 385)

As Silins points out, the defender of EE knowingly objects to Equal Justification. Evidential Externalism is partially constituted by a rejection of is doing better than Barry. Kennedy ([2010]) notes something similar.

something like Equal Justification. Thus, if Silins' argument against EE is really going to have force, a defense of BCNB cannot rely merely on supposing that Equal Justification is true.

With these preliminary comments out of the way, I will now respond to objections to the claim that Gina really is epistemically worse off than Beth.

3.1 Conditional Degrees of Belief

One might think that Gina really isn't justified in giving credence 1 to R , since before getting evidence, Gina is such that $cr(R|SR) = 0.9$. Instead, she would be fully justified in giving 0.9 credence to R . But then the scenario above doesn't show BCNB to be false.

When put this way, the objection simply fails. If R and SR are both evidence, then the value of $cr(R|SR)$ is irrelevant to the final value for $cr_E(R)$.¹¹ Consider an example. Let S = "it is snowing," and TS = "my friend tells me it is snowing." Suppose, as is plausible, that $cr(S|TS) < 1$. Nevertheless, imagine that I see that it is snowing and at the same time hear my friend tell me that it is snowing. Even if $cr(S|TS)$ is low, there is nothing odd about my new credences being such that $cr_E(S) = cr_E(TS) = 1$. The externalist about evidence can say the same thing about the stop sign case above. From the mere fact that $cr(R|SR) = 0.9$, there is nothing objectionable about the fact that $cr_E(R) = 1$, given that Gina has R and SR as evidence.

But one might press on with the objection. In the snow example just given, full credence in S doesn't go via full credence in TS . There are different routes to full belief in each of these propositions: my credence in S is a result of some visual experience and my credence in TS is the result of some auditory experience. If we are to understand the stop sign case in this way, then we must say there are different routes to credence in R and credence in SR . Perhaps there are two different experiences that lead Gina to full credence in each of these propositions. But if that's the case, then there should be similarly different routes to belief in each of these propositions for Gina's internal twin, Beth. If there are, then one might think Beth would be fully justified in giving R full credence, too. So, if Beth is such that $cr_E(R) = 0.9$, $cr_E(SR) = 1$ and Gina is such that $cr_E(R) = cr_E(SR) = 1$, then Beth would be doing worse than Gina. This, then, wouldn't tell against BCNB.

¹¹This is a consequence of the fact that one must update on one's *total* evidence (in Gina's case: R and SR) and not just part of one's evidence.

So, the objection can be restated more forcefully. The objection is that Gina cannot be fully justified in setting $cr_E(R) = 1$ solely as a result of the fact that $cr_E(SR) = 1$ and the fact that $cr(R|SR) = 0.9$. I think there is something plausible in this thought. It does seem inappropriate for my credence in R to be 1, if it was based solely on the fact that $cr_E(SR) = 1$ and $cr(R|SR) = 0.9$. Nevertheless, there is a response to this. The response is to note that the story of Gina and Beth is consistent with Gina's $cr_E(R) = 1$ not being based on her credence in SR or her conditional credence, $cr(R|SR)$. Thus, the story does not violate the plausible idea that it is inappropriate for $cr_E(R) = 1$ if it is solely based on the fact that $cr_E(SR) = 1$ and $cr(R|SR) = 0.9$.

Suppose that Gina and Beth both have some internal experiential state, which is causally responsible for $cr_E(SR) = 1$. Assume that this experiential state fully justifies $cr_E(SR) = 1$ in both Gina and Beth. Nevertheless, we can maintain that this internal experiential state in Gina also justifies full credence in R , though it doesn't justify full credence in R for Beth. Thus, if Gina were to give full credence to R , it would be fully justified, but it would not be fully justified in virtue of the fact that it was based on her credence in SR or her conditional credence, $cr(R|SR)$.

Now, one might object that if the internal experiential state in Gina justifies full credence in R , then this same experiential state must do so for Beth as well. And if that's the case, then Gina and Beth are going to come out as equally justified. But this thought is incorrect. Though Gina and Beth are internal duplicates, this does not mean that their internal states share all the same properties. In particular, if Beth is a brain-in-a-vat, then their internal states will have different causal histories. Unless we are assuming internalism about justification, it is fair to suppose that such properties are relevant to the justificatory work that such internal states can do. To deny this is to simply deny Equal Justification from the start. Thus, this response gives us no reason to reject the counterexample.

3.2 More Ways to Be Wrong

One might have a different objection to the counterexample.¹² One might argue that there's more room to be wrong about how things *are* than about how things *seem*, and so this licenses different attitudes toward SR compared to R . Although Gina bears the appropriate external relation to the truth of both propositions, there are differences between them that motivate the

¹²Thanks to an anonymous referee for alerting me to this response.

different attitudes that Gina has to them. This would show that Gina is doing nothing wrong, and so Gina is not in fact doing worse than Beth.

On one way of understanding the claim that there's more room to be wrong about how things seem, I think it is false. If I'm not very good at detecting how things *seem* to me, but very good at detecting how things *are*, then I might very easily go wrong about propositions like *SR* but not propositions like *R*. Suppose we say that there's more room to be wrong about some set of facts, *F*, compared to another set, *G*, just in case you're worse at detecting facts in *F* compared to facts in *G*. Then it is certainly possible for there to be more room to be wrong about how things seem than about how things are.

There's a different way of understanding the claim that there's more room to be wrong about *R* than about *SR*, however. One could say that there's more room to be wrong about *R* than *SR* because in all worlds where *R* is true, *SR* is true, but there are *SR* worlds where *R* is false. If that were true, then there would be a sense in which there's more room to be wrong about how things are than about how things seem. However, this is clearly false: there are *R* worlds where *SR* is false.

There is a final way of understanding the claim that there's more room to be wrong about how things are compared to how things seem. We could say that there's more room to be wrong about *R* compared to *SR* since it is consistent with how things seem to Gina that *R* is false, but it is not consistent with how things seem to her that *SR* is false. However, this clearly begs the question against the defender of EE, by assuming that consistency with how things seem to the agent is somehow privileged. I conclude that this objection fails.

3.3 Bootstrapping

A final objection to the Gina and Beth case involves bootstrapping phenomena.¹³ Bootstrapping happens when an agent is able to manufacture justification for some proposition, seemingly out of thin air. Consider the classic example of bootstrapping: in looking at my gas gauge I conclude that my tank is full. I also conclude that the gauge reads "Full". If I do this many times, it looks as if I can conclude, just from looking at the gas gauge, that the gas gauge is perfectly reliable. This strikes many as inappropriate epistemic behavior.

This applies to the case of Gina and Beth in the following way. One

¹³See Vogel ([2000]).

might worry that if Gina is such that $cr_E(R) = cr_E(SR) = 1$, then she can illegitimately bootstrap her way to the conclusion that the way things seem to her are perfectly correlated with the way that things are. Above I claimed that Gina would be doing better if $cr_E(R) = cr_E(SR) = 1$. But if having such credences allow her to bootstrap, then perhaps Gina *wouldn't* be doing better if $cr_E(R) = 1$ and $cr_E(SR) = 1$.

Bootstrapping phenomena are known to be difficult for epistemic externalists to handle. Though I think there are plausible things that the externalist in general, and an adherent of EE in particular, can say about bootstrapping worries such as these, I will not consider them here.¹⁴ Instead, I will note two things. First, as will be obvious in the next section, bootstrapping worries arise for the adherent of EI in a similar way. Given this, bootstrapping problems aren't due to commitment to EE alone, and so cannot be used in an argument against EE. Second, a variant of the case of Gina and Beth can be constructed where bootstrapping worries do not arise. Here is the case:

There's a red stop sign in front of Gina. It is also true that it seems to Gina as if there's a red stop sign in front of her. However, Gina is not very reflective, and so Gina has no belief about whether it seems that there is a stop sign in front of her. In virtue of her situation, R is evidence for Gina. Thus, according to EE it should be assigned a credence of 1. Despite this, Gina is such that $cr_E(R) = 0.9$. Beth, on the other hand, is in quite different situation. While it does seem like there's a red stop sign in front of Beth, there isn't a red stop sign in front of her. Since Beth is Gina's internal twin, Beth does not have any beliefs about whether it seems to her that there is a stop sign but is such that $cr_E(R) = 0.9$. Since Beth is in a different situation, and R is not evidence for her, it is completely appropriate for Beth to be such that $cr_E(R) = 0.9$.

In this situation, Beth is doing better than Gina. There is a red stop sign in front of Gina that is responsible for Gina's internal experiential state. This makes R evidence for Gina and justifies full credence in R . Nevertheless,

¹⁴For a general response, see Kornblith ([2009]). He responds to bootstrapping worries on behalf of the reliabilist. His response can be read as a general attempt to disarm our discomfort with subjects that bootstrap, by contrasting those who are reliably accurate bootstrappers and those who are not. If, then, Gina is a sort of reliably accurate bootstrapper then we can disarm the intuitions one might have against such bootstrapping behavior.

Gina doesn't give R full credence. There isn't a stop sign in front of Beth, and nothing like a stop sign is responsible for Beth's internal experiential state. Accordingly, this state does not make R evidence, and Beth doesn't treat it this way. Beth is responding more appropriately to her situation, which is different than Gina's. But Gina is in a good case and Beth is in the bad case. So, BCNB is false. Further, since Gina has no beliefs about how things seem to her, there is no way for her to bootstrap her way to potentially objectionable conclusions.

3.4 Summary

I have argued that the Gina and Beth example raises significant doubts about BCNB. If that's right, then premise 2 of Silins's argument can be rejected, and the argument against EE fails.

But one may not be convinced. One might maintain that BCNB holds up in the Gina and Beth example. If one maintains this it is because one thinks that Gina isn't *really* doing worse than Beth in the situation described. Intuitively, one might claim, Gina doesn't have R as evidence as claimed. That evidence isn't available to Beth, and for all Gina can tell, she's in the same situation that Beth is in. In general, one might point out that there are skeptical scenarios that Gina cannot rule out according to which R is false. Accordingly, Gina is being reasonable in giving R a credence of 0.9, because this credence value reflects a reasonable amount of doubt about her perceptual abilities. Thus, one might conclude, the example above does not cast doubt on BCNB.

In the next section I show that even this kind of response does not save Silins's argument against EE. Instead, such a response actually causes trouble for EI, and shows Silins's argument to be flawed.

4 The Tu Quoque

In this part of the paper I respond to Silins's argument, on the assumption that one is not convinced by the Gina and Beth example. I will argue that if one insists that the BCNB holds up in the stop sign example, then there are arguments, structurally similar to Silins, that push one towards intolerably internal accounts of evidence. To stop this slide to extreme internalism, one must respond to Silins's argument. But the response that must be given can be co-opted by the adherent of EE.

To make this point, I will focus specifically on Silins's formulation of EI, and show how an argument just like his can be used against EI. However, the

discussion is broader than this specific issue. It is relevant and important to any attempt to answer the Evidence Question.

4.1 Doxastic Internalism and Mental Internalism

Silins says that two agents are internal twins when they share the same non-factive mental states. Call such twins *mental twins*. Silins then defends:

mEI: Necessarily, if A and B are mental twins, then A and B have the same evidence.

Consider a different account of evidence, suggested by Howson & Urbach ([1993], p. 99), according to which E is evidence just in case one fully believes E (e.g., $cr(E) = 1$).¹⁵ This is an account of evidence, according to which evidence cannot differ between *doxastic twins*, agents who share all the same doxastic states. We could state the views as follows:

dEI: Necessarily, if A and B are doxastic twins, then A and B have the same evidence.

Now, dEI and mEI are different theses. However, given the statement of the views, it may be hard to find cases where they diverge. For if A and B are mental twins, then they are doxastic twins. So if mEI says that two agents have the same evidence, so does dEI. There are possible situations where A and B are doxastic twins but not mental twins. In these situations, dEI must say that A and B have the same evidence. But mEI need not differ from dEI since mEI takes no stand in such cases. These are simply cases where the antecedent of mEI is not satisfied. However, if mEI and dEI really are different theses, then there must be *some* cases where they diverge. Thus, I'll assume that there are cases where dEI says that two agents have the same evidence, but mEI does not.¹⁶

4.2 An Argument Against mEI

Given this preparatory work, consider an argument against mEI that is structurally identical to the argument Silins gives against EE. Here's how

¹⁵Peter Milne ([2003]) cites this kind of view approvingly: "We do not enquire as to why the agent comes to assign full belief to e in the interval from t_0 to t_1 . This change is, in Howson and Urbach's felicitous phrase, exogenous (1993: 106). Experience may prompt the change, even bring it about causally, perhaps, but the details of the story of how that comes about are of no moment." (p. 283).

¹⁶It is worth noting that Silins makes a similar assumption when discussing EE, and for a similar reason (see section 2).

that goes:

Good Case:

It appears to George as if there is a gorilla walking across the basketball court. However, George doesn't notice that it appears to him this way.¹⁷ It also appears to George that he is told by a neuroscientist that it did appear to him as if there was a gorilla. George does notice this appearance. Accordingly, George is fully confident that he appears to be told this, but isn't fully confident that it did appear to him as if there was a gorilla. Letting AG = "it appears to me that there is a gorilla walking across the basketball court", and ATG = "I appear to be told that it appears to me that there is a gorilla", George's beliefs are like this:

$$cr_E(AG) = 0.9 \quad cr_E(ATG) = 1 \quad cr(AG|ATG) = 0.9$$

Bad Case:

¹⁷One might wonder how it is possible for a person to have an appearance, and yet not notice it. There is mounting evidence that such a scenario is possible. Consider this summary of a study by Simons, et. al.:

...researchers have demonstrated that people fail to notice changes in the information that is visually available to them (Simons 2000). Interestingly, people often cannot describe the change that has taken place, but do demonstrate traces of memory of what they saw before the change. For example, an experimenter holding a basketball stopped pedestrians to ask for directions (Simons et al. 2002). While the pedestrian was giving directions, a group of people (confederates in the experiment) walked between the experimenter and the pedestrian. During this interruption, the experimenter handed the basketball to one person in the group. After giving directions, the pedestrian was asked if he or she noticed any sort of change during the brief exchange with the experimenter. Most did not. However, when led to think about a basketball, the pedestrian did recall seeing it at the beginning of the exchange, and some even recalled specific features of the ball. So, while the participants failed to explicitly notice that a change took place, they did hold accurate implicit memory representations of both the pre- and post-change image. (Chugh & Bazerman, [2007], p. 5)

This provides some evidence that it is possible for it to appear to one as if something is the case (e.g., a basketball was handed off), even if one doesn't notice this. The fact that there are accurate memories of the change imply that it did appear as if something changed, even if the agent didn't consciously notice this. Also see Simons & Chabris ([1999]). For a general philosophical discussion of these cases, see Block ([2008]).

Bertrand is a doxastic twin of George. However, it does not appear to Bertrand as if there is a gorilla walking across the court. It does, though, appear to Bertrand that he is told by a neuroscientist that it did appear to him as if there was a gorilla, and he notices this. Bertrand is thus fully confident that he appears to be told this, but isn't fully confident that it did appear to him as if there was a gorilla. Accordingly, Bertrand's beliefs are like this:

$$cr_E(AG) = 0.9 \quad cr_E(ATG) = 1 \quad cr(AG|ATG) = 0.9$$

Bertrand and George are doxastic twins, so dEI says that they have the same evidence. Thus, neither Bertrand nor George is doing epistemically better than the other according to dEI. However, there are important differences in their mental states. It appears to George as if there is a gorilla, while it does not appear like this to Bertrand. If mEI is to be interestingly distinct from dEI, then mEI will say in cases like this that George has more evidence than Bertrand.

Now, mEI does not give necessary and sufficient conditions for a proposition being evidence. Thus, it doesn't tell us how to figure out what George and Bertrand have as evidence in this situation. However, a plausible way of filling out mEI will say that while only ATG is evidence for Bertrand, both AG and ATG are evidence for George. This is because it appears to George as if there is a gorilla and it appears to George as if he is told this, while it does not appear to Bertrand as if there is a gorilla.¹⁸ An adherent of mEI could, of course, reject that in this case AG and ATG are evidence for George, but as mentioned above, if mEI is to be interestingly different than dEI, there are bound to be cases like this.¹⁹ Thus, there will be cases like this where mEI says that Bertrand is doing epistemically better than George. Bertrand's credences result from conditionalizing on his evidence, whereas George's credences do not.

¹⁸One might worry that I'm assuming that mEI is what Silins calls a *mentalist conception of evidence*. "According to mentalist conceptions of evidence, one's evidence consists only of one's mental states or facts about one's mental states." (p. 394) It is true that this case considers only evidence propositions describing mental states or facts about those mental states (e.g., how one is appeared to). But this is inessential to my argument. For instance, I could have made the relevant propositions G (that there was a gorilla on the court) and TG (that I was told that there was a gorilla), where G is evidence because AG is true of George and TG is evidence because ATG is true. So, this feature of the case is inessential to my argument.

¹⁹Again, it is worth noting that Silins must make a precisely analogous assumption in arguing against EE.

Note, however, that Bertrand is in the bad case compared to George who is in the good case. After all, Bertrand is being deceived by what it appears the neuroscientist is telling him. It appears to Bertrand as if the neuroscientist says that it appears to him as if there is a gorilla. But it *doesn't* appear to him as if there is a gorilla. George, on the other hand, is not being deceived at all. So, we have a situation where mEI says that Bertrand in the bad case is doing better than George in the good case. Thus, the adherent of dEI can criticize Silins's mEI in just the way that Silins criticizes EE: mEI violates BCNB.

The defender of mEI could respond that mEI does not violate the letter of BCNB. This is correct. BCNB was written specifically with the defender of mEI in mind and claims that the agent in the bad case is never epistemically better off than his *mental twin* in the good case. George and Bertrand are not mental twins, they are only doxastic twins. So, BCNB doesn't apply in this case.

This is not a promising response. Recall that this argument is aimed at one who maintains BCNB in light of the Gina and Beth scenario. BCNB is maintained by saying that Gina isn't really doing worse than Beth since Gina doesn't *really* have R as evidence. But *why* maintain that Gina doesn't have R as evidence? One explanation is as follows: for all Gina can tell, she's in Beth's situation, and in Beth's situation R is false. But if that is the response to the Gina and Beth scenario, then surely one should say that George doesn't *really* have AG as evidence, either. After all, for all George can tell, he's in Bertrand's situation, and in Bertrand's situation, AG is false. BCNB would thus generalize to the George and Bertrand scenario.

A different explanation for why BCNB is true in the case of Gina and Beth goes as follows: Gina and Beth have the same perspective on the world, and if they have the same perspective, then they must have the same evidence.²⁰ Notice, however, that there are multiple ways of understanding what it means to have the same perspective. We could understand 'same perspective' as 'same non-factive mental states'. But one could alternatively maintain that 'same perspective' is to be understood as 'same conscious mental states', or as 'same doxastic states'. If we choose this latter way of understanding 'same perspective', then the response to the Gina and Beth case requires one to claim that George and Bertrand have the same evidence and so mEI is false. Again, BCNB generalizes to this case. Of course, the defender of mEI might encourage us to understand 'same perspective' as 'same non-factive mental states'. But without further argument, this is

²⁰Thanks to an anonymous referee for suggesting this explanation.

merely a stipulation that mEI is true.

When initially introduced, BCNB was formulated specifically for the purposes of a defender of mEI. However, any plausible way of defending BCNB in the Gina/Beth case, allows the general idea behind the principle to extend to the case of George and Bertrand. This general idea dictates that George cannot be doing worse than Bertrand.

So, there is an argument against mEI structurally similar to Silins's argument against EE. Importantly, note how a defender of mEI would have to respond to this argument. To respond, the defender of mEI must say that there are situations where an agent has some proposition, P , as evidence because some relation between the agent and P actually obtains, even if the agent is in no position to realize that the relation obtains. Once this move is allowed, the defender of EE can respond to Silins's argument analogously. The mere fact that for mEI these relations are internal to the agent, and for EE they need not be, is irrelevant.

We have seen that mEI and EE are vulnerable to precisely the same kind of argument. If the argument shows EE to be false, then mEI is shown to be false, too. One natural response is to claim that the defender of mEI really meant to be claiming that an agent's evidence is determined by those mental states to which an agent has conscious *access*. Accordingly, this *tu quoque* would be no objection to the spirit of EI.

Note, however, that this is clearly *not* Silins's position. He is clear that he does not take the internal to be exhausted by those things the agent has access to.²¹ Other internalists seem to be in agreement with Silins here.²² So, Silins and others will need to defend mEI against this kind of argument.

This is important. It shows that a plausible version of Evidential Internalism, mEI, is in exactly the same situation as EE with respect to Silins's argument.

5 Accounts Immune to the Argument

In response to this problem, one might suggest that the defender of EI *should* opt for an account of evidence that does not violate BCNB, never mind that Silins does not. Perhaps dEI is such a view, or perhaps there is some other

²¹“Others might use different readings of ‘internal’, and propose narrower supervenience bases for two thinkers to have the same evidence. For example, an internalist might say that ‘internal’ mental states are just those non-factive mental states which are consciously accessible. [...] Evidential Internalism is not what we may call an access thesis in epistemology.” (p. 377).

²²See, for instance, Feldman ([1988], p. 101).

view that does not violate BCNB. In what follows, I will sketch a schematic outline of Silins's argument, and show that at least two accounts of evidence appear to be immune to it. However, I will argue that there are problems with each of these. Thus, we must conclude that Silins's argument is flawed: sometimes the bad case really can be better.

Here is a schematic account of Silins's argument. The argument takes aim at some target account of evidence. The target account says that differences in some supervenience base (SB) can make an evidential difference. We then construct a scenario where there is an SB-difference between agents who are, in some important sense, twins. Perhaps they have the *same perspective*; perhaps things *look the same to them*; perhaps *for all they can tell, they are in the same situation*. We then construct a scenario so that the target account says that the SB-difference makes the good case twin have more evidence than the bad case twin. However, we construct the scenario so that both twins have doxastic states that treat as evidence what the target account says is the *bad case* twin's evidence. Thus, the target account allows a scenario where the bad case twin is epistemically better off than his twin in the good case. This, however, conflicts with the general idea behind BCNB.

First note that one cannot give a Silins-style argument against dEI. To give the argument we need a scenario where the agent in the good case has more evidence than the agent in the bad case, but his doxastic state does not respond to it. But if it is the agent's credence in E that determines whether or not E is evidence, then it is impossible for the agent's doxastic state not to respond to his evidence. So, we can't construct the scenario needed for the argument.

The second kind of account of evidence to consider is a very strong form of evidential internalism. According to this view, there can be no evidential differences between two agents who are consciously aware of the same internal states. Call such a view *Conscious State Evidential Internalism* (cEI). It is plausible that cEI can avoid this kind of argument. This is because it is doubtful that there can be differences in the states two agents are consciously aware of and yet *not* be differences in their perspectives, no matter how we understand what it is to have a certain perspective. If that's right, then we won't get the kind of scenario we need for the argument against cEI. For we would need a situation where there is a difference in the conscious mental states (SB) of two agents, and yet for BCNB to apply, their perspectives would need to be the same. If that's not possible, then the argument is avoided.

We have, then, two versions of EI that seem to avoid the argument:

Conscious State Evidential Internalism (cEI) and Doxastic Evidential Internalism (dEI). How do these fare as accounts of evidence?

Consider dEI first. Take two agents with identical doxastic states. Suppose that both fully believe that they are seeing the sunset over the Grand Tetons. One of the twins really is having an experience as of a sunset over the Grand Tetons, while the other is not. Since they have identical doxastic states, there can be no epistemic differences between them in virtue of their distribution of beliefs. According to dEI, these doxastic twins must have the same evidence. Thus, there can be no epistemic differences between them in virtue of the evidence that they have. But this does not allow us to distinguish the two in the way that we would like. There clearly is an epistemic distinction to be drawn between the two agents.

Consider, then, cEI. This account is also problematic. According to cEI there is a *very* narrow supervenience base for our evidence: those states of which we are consciously aware. But it is unlikely that such a narrow supervenience base can really do all the justificatory work that we think our evidence does. Just think of all the information we have stored in memory, that seems to do justificatory work, and yet only a small portion of which we are consciously aware at any time. I think I am justified in having a high degree of belief about things like my name and where my office is. I think I am justified in such degrees of belief even when I am not consciously reflecting on any evidence relevant to such matters, but I think I am justified precisely because I have evidence that is relevant to such matters. This very basic fact is not captured by cEI. In sum, cEI presents us with a very restrictive conception of evidence. It isn't clear that this conception of evidence is always—or even usually—the conception of evidence that is worth caring about.

There is, after all, a more objective sense of 'evidence' where what evidence an agent has does not depend on subtle details about what happens to be conscious to the agent with whom we are concerned. It is very plausible that it is this more objective sense of 'evidence' that is at issue when we talk about rational degrees of confidence, or offer formal Bayesian models of belief change. For example, the relevant evidence in a legal trial does not depend solely on how things consciously seem to members of the jury. The evidence in a scientific setting surely doesn't solely depend on how things consciously appear to the scientists. Certainly this is true if we think that regular scientific theories and court verdicts are well-supported by the evidence. Conscious appearance may play a role in the story, and it may be an important one, but it need not be the *whole* story.

In fact, Silins's own version of Evidential Internalism (mEI) is one ac-

count of evidence according to which evidence is more objective, and not solely dependent on conscious experience. But once we notice that there is this sense of ‘evidence’ not so intimately tied to what is consciously entertained, then an argument like Silins’s against EE is unconvincing. It is unconvincing because the defender of any account of evidence—including mEI—will have to answer an exactly similar argument to stop the slide to the kind of extreme-internalism about evidence represented by cEI and dEI.

6 Conclusion

Near the end of his article, Silins writes:

The point I want to emphasize is that, if Evidential Externalism is true, then counterintuitive assessments of beliefs in the good case turn out to be correct. Intuitively, our ordinary partial beliefs are just fine as they are: we are neither more nor less confident in our ordinary beliefs than we should be. If Evidential Externalism is true, however, we are sometimes less confident than we should be, given our evidence. (p. 390)

There is some truth in this statement. EE does lead to assessments of belief that can be counterintuitive in some instances. But I hope to have shown that this is a shared problem, one that many other accounts of evidence face as well.

There is an important general point to notice here: requiring that evidence supervene on mental states, as mEI does, does nothing to insure that the bad case won’t sometimes be better. This is because our thoughts about whether the bad case can be better have to do with intuitions about whether or not an agent can *tell* that a proposition is evidence, and not with supervenience. It is natural to think that if two agents are in the same situation, for all they can tell, then the bad case certainly couldn’t be better. However, being able to tell that something is evidence is one thing, and having evidence supervene on the mental is a very different thing. Thus, accounts of evidence that require evidence to supervene on the mental will lead to counterintuitive assessments of belief in much the same way as those that do not require evidence to supervene on the mental. So the cases Silins directs us towards couldn’t possibly provide a reason in favor of EI against EE.

Recognizing this is important if we hope to answer the Evidence Question. It seems that we think of evidence as the kind of thing that is supposed to play two very different sorts of roles. On the one hand, evidence is supposed to be the kind of thing connected with truth, and which helps to justify

beliefs. Given this, when we acquire evidence, it should directly affect our beliefs in some way, mandating certain changes in belief in accordance with how the world is. On the other hand, evidence is supposed to be something that the agent can use, from his own perspective, for epistemic guidance. These two desiderata, however, tend to pull in opposite directions. I hope to have shown that this tension is not just a tension faced by one who thinks about evidence more externally.

7 Appendix

In the body of the paper, I go along with Silins in assuming that one's evidence receives credence 1. Here I note that the assumption that evidence receives credence 1 is tangential to my arguments.

According to some, a better picture of evidence allows for graded evidence: evidence propositions that receive credence values less than 1.²³ Strictly speaking, if one countenances graded evidence there are not evidence *propositions*, but rather an evidence *partition*, with different weights assigned to different parts of the partition. I will show that thinking of evidence in this way does not affect the arguments.

As noted in section 5, a Silins-style argument works so long as the agent in the good case has more evidence than his twin in the bad case, but the agent in the good case doesn't have a doxastic state that responds to his greater amount of evidence. This generates the result that the agent in the good case is epistemically worse off than his twin in the bad case.

If all evidence receives credence 1, then an agent has more evidence than his twin just in case his evidence set contains more propositions. This makes it simple to set up and think about the scenario needed to generate the argument. Things are not so simple if we countenance graded evidence. However, there are still things we can say about when an agent has more evidence than his counterpart. For instance, suppose that two agents have the same evidence partition. In this case we can say that an agent has more *E*-evidence than his counterpart just in case *E* is a member of both their evidence partitions, but the one twin is justified in believing *E* to a higher degree than his counterpart. So long as EE and mEI both say that sometimes agents in the good case have more evidence than their twins in the bad case, in *this* sense, we can construct a Silins-style argument when countenancing graded evidence.

²³For discussion of this issue, see, for instance, Jeffrey ([1983]), Field ([1978]), Garber ([1980]), Christensen ([1992]), Joyce ([2004]), and Weisberg ([2009]).

For example, consider Gary and Barry. Suppose that Barry's evidence partition is:²⁴

$$E_B: \{\langle B \wedge SB, 0.89 \rangle, \langle B \wedge \neg SB, 0.01 \rangle, \langle \neg B \wedge SB, 0.09 \rangle, \langle \neg B \wedge \neg SB, 0.01 \rangle\}$$

If we adopt Jeffrey Conditionalization²⁵ as the correct way to update on graded evidence, this results in $cr_{E_B}(B) = 0.9$ and $cr_{E_B}(SB) = 0.98$.

Suppose that Gary's evidence partition is:

$$E_G: \{\langle B \wedge SB, 0.97 \rangle, \langle B \wedge \neg SB, 0.01 \rangle, \langle \neg B \wedge SB, 0.01 \rangle, \langle \neg B \wedge \neg SB, 0.01 \rangle\}$$

This results in $cr_{E_G}(B) = 0.98$ and $cr_{E_G}(SB) = 0.98$. In such a case, Gary has more B -evidence than Barry. In the original case, EE says that Gary has B as evidence whereas Barry does not. In much the same way, and for the same reason, EE could say that Gary has more B -evidence than Barry.

Suppose, then, that Gary's evidence is E_G and Barry's evidence is E_B . Suppose, however, that Gary is such that $cr_{E_G}(B) = 0.9$ and $cr_{E_G}(SB) = 0.98$. Here we would get the result that Barry, in the bad case, is epistemically better off than Gary in the good case. Barry's credences result from Jeffrey Conditionalizing on Barry's evidence, but Gary's do not result from Jeffrey Conditionalizing on Gary's evidence. Clearly, something analogous could be constructed for the case of George and Bertrand.

In the example above, Gary and Barry both have the same evidence partition, but with different weights. It is also possible that Gary has more evidence than Barry by having an evidence partition that is more fine-grained (intuitively, one that contains *more* propositions). For instance, EE might imply that Gary's evidence is E_G , whereas Barry's evidence is:

$$E_{B*}: \{\langle SB, 0.98 \rangle, \langle \neg SB, 0.02 \rangle\}$$

As in the original case, suppose that for both Barry and Gary, $cr(B|SB) = 0.9$. Jeffrey Conditionalization tells us that if Barry updates on his evidence, his credence in B should be:

$$cr_{E_{B*}}(B) = cr(B|SB) \times cr_{E_{B*}}(SB) + cr(B|\neg SB) \times cr_{E_{B*}}(\neg SB)$$

²⁴Note that I use ordered pairs of the form $\langle P, n \rangle$ to indicate that P is evidence to degree n .

²⁵See Jeffrey ([1983], ch. 11).

From this it follows that $cr_{E_{B^*}}(B) < 0.902$. Suppose, however, that Gary and Barry have identical doxastic states where their credence in B is less than 0.902. Then it would follow that Barry's credences result from Jeffrey Conditionalizing on Barry's evidence, but Gary's do not result from Jeffrey Conditionalizing on Gary's evidence. Again, we could clearly construct something analogous for the case of George and Bertrand.

What is crucial here is that both EE and mEI sometimes have the consequence that the agent in the good case has more evidence than the agent in the bad case, but that the agent in the good case has a doxastic state that does not respond to that extra evidence. As we've seen, moving to graded evidence does not prohibit this situation from obtaining. It simply makes it more complicated to say just what it is to have more evidence.

Thus, so long as there is something objectionable about the bad case being better, moving to graded evidence does not allow one to evade a Silins-style argument. It seems to me that if one finds such a situation objectionable when evidence receives credence 1, then one will find it objectionable when evidence is graded. But suppose one finds this situation less objectionable with graded evidence. Then, perhaps, moving to a framework with graded evidence provides a way to respond to my argument against mEI. But if so, then Silins's argument against EE can be evaded in just the same way. Thus, the assumption that evidence receives credence 1 in no way plays a role in the arguments above. It simply makes them more easy to state.²⁶

References

- [1] Block, N. [2008]: 'Consciousness and Cognitive Access', *Proceedings of the Aristotelian Society*, **108**, Part 3, pp. 289-317.
- [2] Christensen, D. [1992]: 'Confirmational Holism and Bayesian Epistemology', *Philosophy of Science*, **59**, pp. 540-57.
- [3] Chugh, D. & Bazerman, M. [2007]: 'Bounded Awareness: What You Fail to See Can Hurt You', *Mind and Society*, **6**, pp. 1-18.
- [4] Feldman, R. [1988]: 'Having Evidence', in David Austin (ed.), *Philosophical Analysis*, Dordrecht: Kluwer, pp. 83-104.

²⁶Acknowledgements: I am grateful to Hilary Kornblith for very helpful conversation on this material. I would also like to thank an anonymous referee from this journal for helpful comments.

- [5] Field, H. [1978]: 'A Note On Jeffrey Conditionalization', *Philosophy of Science*, **45**, pp. 361-367.
- [6] Fitelson, B. [2001]: *Studies in Bayesian Confirmation Theory*, Doctoral Dissertation, University of Wisconsin, Madison.
- [7] Garber, D. [1980]: 'Field and Jeffrey Conditionalization', *Philosophy of Science*, **47**, pp. 142-145.
- [8] Howson, C. & Urbach, P. [1993]: *Scientific Reasoning: The Bayesian Approach (2nd Edition)*, Peru, IL: Open Court.
- [9] Jeffrey, R. [1983]: *The Logic of Decision, 2nd Edition*, Chicago: The University of Chicago Press.
- [10] Joyce, J. [2004]: 'Williamson on Evidence and Knowledge', *Philosophical Books*, **45**, pp. 296-305.
- [11] Kennedy, M. [2010]: 'Naive Realism and Experiential Evidence', *Proceedings of the Aristotelian Society*, **90**, pp. 77-109.
- [12] Kornblith, H. [2009]: 'A reliabilist solution to the problem of promiscuous bootstrapping', *Analysis*, **69**, pp. 263-267.
- [13] Maher, P. [1996]: 'Subjective and Objective Confirmation', *Philosophy of Science*, **63**, pp. 149-174.
- [14] Neta, R. [2008]: 'What Evidence Do You Have?', *British Journal for the Philosophy of Science*, **59**, pp. 89-119.
- [15] Silins, N. [2005]: 'Deception and Evidence', *Philosophical Perspectives*, **19**, pp. 375-404.
- [16] Simons, D. J., & Chabris, C. F. [1999]: 'Gorillas in our midst: Sustained inattentive blindness for dynamic events', *Perception*, **28**, pp. 1059-1074.
- [17] Simons, D. J. [2000]: 'Current approaches to change blindness', *Visual Cognition*, **7**, pp. 1-15.
- [18] Simons, D. J., Chabris, C. F., Schnur, T., Levin, D. T. [2002]: 'Evidence for preserved representations in change blindness', *Consciousness and Cognition*, **11**, pp. 78-97.

- [19] Vogel, J. [2000]: ‘Reliabilism Leveled’, *The Journal of Philosophy*, **97**, pp. 602-623.
- [20] Weisberg, J. [2009]: ‘Commutativity or Holism: A Dilemma for Conditionalizers’, *British Journal for the Philosophy of Science*, **60**, pp. 793-812.
- [21] Williamson, T. [2000]: *Knowledge and Its Limits*, Oxford: Oxford University Press.