

Evidentialist Anti-Skepticism

Evidentialism is the view that knowledge requires evidence, and that the possession of evidence gives us knowledge provided our evidence eliminates all reasonable doubt, and provided we are not in a Gettier-type situation. I will discuss what we should say in response to skepticism if we endorse Evidentialism thus understood. One of the most prominent skeptical arguments involves the premise that one does not know one is not a BIV. This argument goes as follows:

The BIV Argument

- (1) $KH \rightarrow K\sim BIV$
- (2) $\sim K\sim BIV$
- (3) $\sim KH$

The first premise can be seen as an implicit appeal to *closure*. If the underlying assumption is made explicit, the argument goes as follows:

The BIV Closure Argument

- (1) $[KH \ \& \ K(H \rightarrow \sim BIV)] \rightarrow K\sim BIV$
- (2) $K(H \rightarrow \sim BIV)$
- (3) $\sim K\sim BIV$
- (4) $\sim KH$

Many contemporary epistemologists consider this argument an intellectual challenge of the highest magnitude. To deal with it, they either recommend to abandon closure or to introduce a non-epistemic entitlement to dismiss the BIV hypothesis out of hand. Those who think that rebutting the BIV argument calls for such drastic measures are united in endorsing one of the following two theses:

The No-Evidence-At-All Thesis

One has no evidence at all for thinking that one is *not* a BIV.

Generalized:

If according to a skeptical hypothesis, H, one's evidence is the same as it is now, then one has no evidence at all *against* H.

The No-Evidence-That's-Good-Enough-for-Knowledge Thesis

One has some evidence for thinking one is *not* a BIV, but that evidence isn't good enough for knowledge.

Generalized:

If according to a skeptical hypothesis, H, one's evidence is the same as it is now, then one has no evidence that's good enough for knowing $\sim H$.

The rationale for the No-Evidence-At-All thesis is obvious: A person who undergoes envatment does not undergo any change of her evidence. So after envatment, one has exactly the same

evidence one has before envatment. Consequently, one cannot have any evidence for thinking one is not a BIV.

Those who favor the Not-Good-Enough Thesis think we are in possession of some weak non-deductive reasons for thinking we are not envatted. For example, it would appear the technology for keeping a brain alive for an extended period does not yet exist. However, advocates of the Not-Good-Enough Thesis would hasten to add that such evidence doesn't do what evidence must do if it is to give us knowledge: protect us in a robust way against error. It does not so protect us because if we were BIVs we would still think that the technology needed for envatment is not available. The same holds for any evidential item that could be brought to bear against the BIV hypothesis. Therefore, whatever anti-BIV evidence we might have, it isn't good enough for us to know that the BIV hypothesis is false.

Both of these theses strike me as false. For an assessment of the plausibility of them, I recommend considering a different type of skeptical scenario, one in which one is not envatted or otherwise deceived by a mad scientist or some sinister state or rogue agency, but rather by the Easter Bunny. If one is a BIV, one thinks one has hands while in fact one does not. The same is going on if the Easter Bunny is the agent of deception. Being handless, and indeed being altogether disembodied, the Easter Bunny's victim labors under the illusion of having a normal life. The Easter Bunny deception analog to the BIV argument, then, runs as follows:

The EB Deception Argument

- (1) $[KH \ \& \ K(H \rightarrow \sim EBD)] \rightarrow K\sim EBD$
- (2) $K(H \rightarrow \sim EBD)$
- (3) $\sim K\sim EBD$
- (4) $\sim KH$

How would one want to respond to this argument? I'm inclined to think most people take themselves to know that the Easter Bunny does not exist. I, in any case, take myself to know this. Assuming, then, we all think we know the Easter Bunny does not exist, we might wonder on the basis of what evidence we have acquired this bit of knowledge. Let me just mention one evidential item among many: having paws, bunnies lack the manual dexterity required for distributing eggs on a global scale (or any scale, for that matter). Of course, if we think that knowledge requires truth-entailing evidence, then the paw argument is no good at all. But, obviously, if we take ourselves to know that the Easter Bunny does not exist, we don't think that knowledge requires truth-entailing evidence – which is to say we are endorsing fallibilism.

According to **fallibilist evidentialism** as I understand it here, if evidence is to give us knowledge, it need not be truth-entailing, but it must eliminate all reasonable doubt. I take it our evidence leaves no reasonable doubt as to the Easter Bunny's non-existence. That's why, according to fallibilist evidentialism, we know that the Easter Bunny does not exist.

Let's suppose, then, we agree that the following thesis is true:

The EB Non-Existence Thesis

We have evidence for EB nonexistence that's good enough for knowledge.

What's interesting about this thesis is this: if you think it is true, and if you think your body of evidence can be enlarged using deduction (that is, if you accept *closure*), then you should reject the two theses we considered a moment ago—the No-Evidence-At-All and the No-Evidence-That's-Good-Enough theses. Consider first the general form of the No-Evidence-At-All thesis:

The No-Evidence-At-All Thesis

If, according to a skeptical hypothesis, H, one's evidence is the same as it is now, one has no evidence at all against H.

Here is why this thesis is false. In a world in which you are deceived by the Easter Bunny, your evidence is the same as it is in this world. In *this* world, you have excellent evidence for thinking that the Easter Bunny does not exist. Hence, in a world in which you *are* deceived by the Easter Bunny, you *also* have excellent evidence for thinking that the Easter Bunny does not exist. Next, imagine yourself the Easter Bunny's hapless victim. It occurs to you that you can work on your evidence using deduction. You reason as follows:

The EB Closure Step

The EB does not exist.

If so, I'm not deceived by the EB.

Therefore:

I'm not deceived by the EB.

And so it turns out that even if you *were* deceived by the Easter Bunny, you would have an excellent argument for believing that you are *not* deceived by the Easter Bunny. So it looks like we may put forward the following thesis:

The We-Have-Evidence Thesis

We have excellent evidence for thinking that the Easter Bunny deception hypothesis is false even though, if that hypothesis were true, our evidence would be exactly the same as it is now.¹

¹ This thesis can be supported as follows: Ask your self which attitude you should take toward the proposition 'BIVs don't exist'. If it is that of acceptance, you must think there is no reasonable doubt as to the nonexistence of BIVs. For if you thought there was such a reasonable doubt, surely the right attitude would be that of suspension of judgment.

The We-Have-Evidence Thesis strikes me as true. It seems to me, therefore, that the Non-Evidence-At-All Thesis is false.

Next, consider the No-Evidence-That's-Good-Enough Thesis, which in its general form says the following:

The-No Evidence-That's-Good-Enough-for-Knowledge Thesis

If according to a skeptical hypothesis, H, one's evidence is the same as it is now, one has no evidence that's good enough for knowing \sim H.

Let's assume we are in agreement about the following: we *know* the Easter Bunny doesn't exist. Now each of us can perform a little deduction: If the Easter Bunny doesn't exist, one is not deceived by the Easter Bunny. This looks like a rather safe step. It is not a complex and lengthy piece of reasoning that weakens the evidence to which it is applied. So if our anti-Easter-Bunny-existence evidence gives us knowledge of the Easter Bunny's non-existence, one should think that, performing the little deduction just mentioned, that very same evidence puts us in a position to know we are not deceived by the Easter Bunny. Let's put this point in the form of another thesis:

The We-Know-It Thesis

We know that the EB Deception Hypothesis is false even though, if that hypothesis were true, our evidence would be exactly the same as it is now.

If the We-Know-It Thesis is true, then the No-Evidence-That's-Good-Enough for Knowledge Thesis is false. Since the former thesis is exceedingly plausible, it seems to me it supplies us with a good reason to reject the latter thesis.

The arguments for the We-Have-Evidence and the We-Know-It theses rely on *closure*. However, closure is not uncontroversial. Some people reject it. What's nice about the deceptive Easter Bunny and other deceivers of the same ilk is the following: They illustrate just how implausible it is to deny closure, at least when it comes to obviously non-existing agents of deception. If you want to reject my line of reasoning on the ground that I'm making an illegitimate appeal to closure, then you you'd have to say things like the following:

Closure Denial Costs

- You know that the Easter Bunny does not exist, but you don't know that you are not deceived by the Easter Bunny.
- You know that Napoleon is dead, but you don't know that you are not deceived by Napoleon.
- You know that dinosaurs are extinct, but you don't know that you are not deceived by some dinosaur.

As far as abominable conjunctions go, these are particularly egregious. It seems to me we should

avoid them and agree that, if one knows the Easter Bunny doesn't exist, then one also knows one is not deceived by the Easter Bunny. But then the following, rather effective response to the Easter Bunny Deception Argument becomes available to us:

The Easter Bunny Deception Counter-Argument

- (1) $K\sim EBE$
- (2) $K(\sim EBE \rightarrow \sim EBD)$
- (3) $[K\sim EBE \ \& \ K(\sim EBE \rightarrow \sim EBD)] \rightarrow K\sim EBD$
- (4) $K\sim EBD$

Let me briefly sum up: This Easter Bunny non-existence response to the Easter Bunny deception argument is based on *fallibilist evidentialism*. According to this view, the standard we must meet to know is high, but not excessively high: we must be in possession of evidence that eliminates all reasonable doubt. There is no reasonable doubt about the Easter Bunny's nonexistence. Nor is there reasonable doubt about the relevant entailment. Hence one knows that the Easter Bunny deception hypothesis is false.

We have now reached the end of our detour through the territory of Easter Bunny skepticism. Next, let's see whether BIV skepticism can be dealt with in an analogous fashion. So let's consider the following reply to the BIV Argument:

The BIV Counter-Argument

- (1) $K\sim BIVE$
- (2) $K(\sim BIVE \rightarrow \sim BIV)$
- (3) $[K\sim BIVE \ \& \ K(\sim BIVE \rightarrow \sim BIV)] \rightarrow K\sim BIV$
- (4) $K\sim BIV$

To begin with, the second premise seems rather innocuous. If BIVs don't exist, then I'm not a BIV. That's certainly beyond any reasonable doubt. So let's move on to the first premise. Is there any reasonable doubt as to the non-existence of BIVs? Let's dwell for a moment on the relevant evidence. Items that come to mind are the following:

1. Textbooks of neurophysiology don't have a chapter entitled 'Envatment'.
2. Departments of neurophysiology don't offer courses entitled 'Envatment 101'.
3. If you bother to call a renowned neurophysiologist or brain surgeon and ask whether envatment is possible, the answer is going to be 'no'.
4. Essay collections for courses on applied ethics don't have a chapter entitled 'The ethics of envatment'.
5. No known episode of *60 Minutes* has ever investigated let alone asserted the existence of BIVs.
6. There is no known case of someone ever having been sued for or found guilty of envatting a person.

This list could go on. If one puts enough effort into it, it could fill pages. Now, obviously even

the collective force of such a list does not *entail* the non-existence of BIVs. But surely it eliminates all reasonable doubt about the non-existence of BIVs. So, arguably, the proposition that there are no BIVs meets the standard of knowledge that fallibilist evidentialism imposes. We know that BIVs don't exist, and hence, performing an easy deduction, we can know that we are not victims of envatment.²

Unfortunately, if that response to BIV skepticism were to work without a major glitch, it would be too good to be true. So there has to be a major glitch. And there is one. It's called circularity, or question-begging. In response to the claim that I know BIVs don't exist, the skeptics will of course apply the BIV hypothesis to a different target proposition. Instead of saying I don't know I have hands because I don't know I'm not a BIV, they will say I don't know BIVs don't exist because I don't know I'm not a BIV. So they will put forward the following argument:

The BIV Argument #2

- (1) $K \sim BIV \rightarrow K \sim BIV$
- (2) $\sim K \sim BIV$
- (3) $\sim K \sim BIV$

Suppose, then, I respond to this argument with the following counter-argument:

The BIV Counter-Argument #2

- (1) $K \sim BIV \rightarrow K \sim BIV$
- (2) $K \sim BIV$
- (3) $K \sim BIV$

A case could be made that this argument is question-begging. Appealing to the possibility of my being a BIV, the skeptical argument calls into question whether I have knowledge of BIV non-existence. The anti-skeptical reply appeals to my knowledge of BIV non-existence, and thus appeals to the very thing the skeptical argument questions. Skeptics could argue that, in making this anti-skeptical reply, I'm like a desert traveler who reasons he is not hallucinating an oasis because he is having a perceptual experience of an oasis. Since the hallucination hypothesis calls

² Of course this could change. Perhaps at some point in the future, envatment will be a wide-spread and much reported phenomenon. Then it's going to be harder to know one is not a BIV. We are not there yet, though. It's like the difference between living in St. Cloud and living in a crime infested metropolis. In St. Cloud, I know my car is where I parked it. In a crime infested metropolis, I might not. Likewise, in this world we know we are not envatted. In a world in which envatment is common place, it might be difficult to know this.

Let me sum up. I think the best evidentialist response to BIV skepticism goes like this. We have excellent evidence for thinking that BIVs don't exist. This evidence would eliminate all reasonable doubt about the existence of BIVs even if one actually were a BIV. On the basis of such evidence, we know that BIVs don't exist. Hence, performing an easy deduction, we know we are not BIVs.

into question the veridicality of the perceptual experience, rejecting that hypothesis by appealing to the perceptual experience seems illegitimate. Likewise, the BIV hypothesis calls into question, at least by way of implication, the non-existence of BIVs. Therefore, according to the skeptics, appealing to BIV non-existence is an illegitimate way of dismissing the BIV hypothesis.

To assess how serious a problem this is, let's again consider Easter Bunny skepticism:

The Easter Bunny Deception Argument #2

- (1) $K \sim EBE \rightarrow K \sim EBD$
- (2) $\sim K \sim EBD$
- (3) $\sim K \sim EBE$

The charge of question-begging is likely to cause epistemic anxiety in the context of rebutting BIV skepticism. But when it comes to rebutting Easter Bunny skepticism, the charge of question-begging seems *less* of a threat. Consider the short version of the response to the Easter Bunny Deception Argument:

The Easter Bunny Deception Counter-Argument #2

- (1) $K \sim EBE \rightarrow K \sim EBD$
- (2) $K \sim EBE$
- (3) $K \sim EBD$

In a purely *structural* sense, this response is question-begging. It uses as a premise precisely what the second Easter Bunny Deception Argument calls into question: knowledge of Easter Bunny Non-Existence. However, I'm inclined to think that we should distinguish between *structural* and *substantive* question-begging. If an argument begs the question in a purely structural, non-substantive way, it need not be a dialectical defect of that argument.

What, then, makes an anti-skeptical argument question-begging in the substantive sense? An anti-skeptical argument is substantively question begging if (i) it's structurally question-begging and (ii) the skeptical hypothesis under consideration is itself *backed up by good evidence*. For example, if there's a good likelihood the desert traveler is hallucinating, he substantively begs the question if he argues he's not hallucinating because he is after all having an experience of an actual oasis. Likewise, if we had good reasons to suspect the Easter Bunny might actually exist and be in possession of highly effective deception technology, then it would be substantively question-begging to dismiss the Easter Bunny Deception Hypothesis on the ground that the Easter Bunny does not exist. However, in the actual world, the Easter Bunny Deception Hypothesis has zero credibility. Hence the Easter Bunny Deception Counter-Argument is guilty at most of purely structural question-begging.

Perhaps purely structural question-begging is no big deal. Evidentialists could say that our evidence against Easter Bunny existence does *double duty*. It tells us two things at once, namely (i) that the Easter Bunny does not exist, *and* (ii) that we are not deceived by the Easter Bunny. Likewise, evidentialists could say that our evidence against the existence of BIVs does double duty: it tells us (i) that BIVs don't exist, *and* (ii) that we not BIVs. If that's right, then the evidentialist non-existence response to BIV skepticism remains effective even though it is, in a purely structural sense, question-begging.