

Evidentialist Anti-Skepticism

1. The BIV Argument and How One Might Respond to It.

Epistemologists worry about not knowing they have hands. The worry arises from skeptical arguments such as the notorious Brain-in-the-Vat Argument, which goes as follows:¹

The BIV Argument

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

The first premise can be seen as an implicit appeal to the closure principle, *closure* for short.³ 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) $\sim KH$

When epistemologists worry that the conclusion of this argument might be true, their worries are fueled by the undeniable plausibility of the premises. The first premise is an instance of closure, which is a principle with a degree of plausibility so high that some consider it axiomatic. The second premise simply states a logical consequence of the BIV hypothesis and is thus completely safe. And the third premise exploits the undeniable difficulties of making a case against a hypothesis that's deliberately designed to rob one of all possible evidence against it. So it looks like the BIV argument does give epistemologists plenty to worry about.

To rebut the argument, which responses are there to choose from? I will discuss the following three:

1. Non-epistemic entitlements, or knowledge without reasons. According to this response, it's possible to know that skeptical hypotheses are false without having

¹ If you are a brain in a vat, your brain was removed from your skull and is kept alive, floating in a vat. The nerve endings of your brain are stimulated in such a way that you have exactly the sort of experiences you would have if you had a normal body and were enjoying a normal life.

² KH =I know I have hands, $K\sim BIV$ =I know I'm not a BIV.

³ According to the closure principle, we know what our knowledge entails if the entailments themselves are known. So if I know P, and I know that P entails Q, then I know Q. There are problems calling for a more cautious articulation of the principle, but these problems may be ignored in the present context. For a defense of closure and a discussion of these problems, see Hawthorne's "The Case for Closure" in Steup, Matthias and Sosa, Ernest (eds). 2005. *Contemporary Debates in Epistemology*. Malden (MA): Blackwell.

evidence or reasons against them.⁴

2. Closure denial. This response is associated with either Dretske's conclusive-reasons theory or Nozick's sensitivity-based account of knowledge.
3. Evidentialist anti-skepticism. According to this option, appearance notwithstanding, we actually have evidence against the BIV hypothesis, and this evidence is strong enough for knowing that the hypothesis is false.

Of these responses, I think the third is the best. In what follows, I'll explain why.

I won't have much to say about the non-epistemic entitlement response. According to it, knowledge requires reasons except when it comes to knowing that skeptical hypotheses are false. So if I am to know I have hands, I have to have good reasons for thinking I have hands. But when it comes to knowing that I'm not a BIV, then I can enjoy such knowledge even though I have no reason at all for thinking that the BIV hypothesis is false. One problem with this approach is that it makes such knowledge mysterious. If knowledge requires reasons except when it comes to knowing that skeptical hypotheses are false, how can my knowledge of not being a BIV be explained? Another problem is arbitrariness. According to the non-epistemic entitlement response, knowledge requires reasons except when it comes to knowing the falsehood of skeptical hypotheses. What justifies making this exception? What we would need for this response to be more attractive is a principled account of when reasons are necessary for knowledge and when they are not. In the absence of such an account, the non-epistemic entitlement suggestion must be rejected as arbitrary.

2. Dretske's Denial of Closure Response

Next, let's consider the response of denying closure, championed by Fred Dretske. According to Dretske, your belief that P amounts to knowledge if, and only if, you have a conclusive reason for believing P. And your reason for P is conclusive just in case, if P were false, you wouldn't have that reason. Your hand-like experiences give you a conclusive reason for believing that you have hands, for if you didn't have hands you wouldn't have those experiences. However, your hand-like experiences don't give you a conclusive reason for thinking you are not a BIV. They do not because if you were a BIV, you would still have those experiences. The same applies to

⁴ See Stine, Gail. 1976. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies* 29, pp. 249-61 and various papers by Crispin Wright.

any other reason you might have for thinking you are not a BIV. The BIV hypothesis is designed to ensure that, no matter what anti-BIV reason we consider, you would still have it even if you were a BIV. So reasons for thinking you are not a BIV cannot be conclusive. And thus we get the outcome that, although you know you have hands, and you know that your having hands entails your not being a BIV, you don't know that you are not a BIV.

One familiar objection to this response points out that it engenders an abominable conjunction: I know I have hands, but I don't know that I'm not a handless BIV.⁵ Later on I'll mention a few more abominable conjunctions. For now, I want to press a different objection. To me, it seems a main flaw with Dretske's approach is that it restricts the scope of knowledge too much. First of all, as a response to skepticism, Dretske's theory is quite concessive. It concedes that we cannot know skeptical hypotheses to be false. The claimed benefit of this concession is that ordinary knowledge—such as knowledge of one's hands—is preserved. But, on Dretske's theory, how much of ordinary knowledge is really safe? For example, do I now know, according to Dretske's theory, that there is a desk in my office? Well, I'm now here, not in my office. Hence my reason for believing there's a desk in my office is not a present desk-like-experience but rather memories of past desk-like experiences plus a bit of background knowledge: desks are rarely removed from faculty offices. But that reason—the combination of remembered perceptual experiences and general knowledge of desk-removal on campuses—is not conclusive: had my desk been removed or stolen, I would still have that reason. What, according to Dretske's theory, I do know is that *last time I looked* there *was* a desk in my office. But I do not know that there is *now* a desk in my office. It seems to me, though, that this is something I know. That's why I think Dretske's theory is too restrictive.

Here's something else I think I know, but don't know according to Dretske's theory. I think I now know that there is not now an atomic bomb in my basement. My reason for believing that there isn't an atomic bomb in my basement consists of a set of premises about atomic bombs: They are very difficult to acquire. People who have them legally keep them within the confines of military installations. People who have them illegally are unlikely to put them in my basement. But of course these reasons are not conclusive. If due to fantastic coincidence, an atomic bomb had found its way into my basement, I would still have these reasons. So Dretske's theory implies that I do not know that there is not now an atomic bomb in my basement.

⁵ See DeRose, Keith. 1995. "Solving the Skeptical Problem." *The Philosophical Review* 104, pp. 1-52.

According to Dretske, if I want to know whether or not there is an atomic bomb in my basement, I would have to go and look. Until then, I remain in a state of ignorance. Once again, it seems to me Dretske's account is too restrictive.

Like Dretske, I would say that having knowledge requires having a reason. But unlike Dretske, I think that for a reason to give you knowledge, it need not be conclusive. That doesn't mean any old reason can give you knowledge. There are reasons that, while they make believing P more reasonable than not believing P, are not strong enough to give you knowledge of P. My reasons for thinking that I will be alive one year from now are an example of that. How strong, then, must a reason for P be if that reason is to give one knowledge of P? My view is that a reason is good enough for knowledge if it eliminates all reasonable doubt. My reasons for thinking that there is a desk in my office, and that there is not an atomic bomb in my basement, eliminate, it seems to me, all reasonable doubt. That's why I think I know these things.

Dretske thinks that knowledge of P always results from indicators that carry the information that P is true. Indicators are perceptual experiences, memories, as well as various forms of testimony: newspaper articles, books, and what people tell me. I would of course agree that such indicators can give us knowledge. The question is whether they are the only things that can give us knowledge. It seems to me in addition to such indicators, there is something else that can give us knowledge: reasoning applied to propositions expressing bits of common knowledge. But as the two examples I mentioned show, such reasoning need not be conclusive. Dretske would therefore deny that it can give us knowledge; I would insist that it can.

In response to my view, I think Dretske would reply that, even if I'm that liberal about what it takes to know, I will still not be able to rebut the BIV Argument. For no matter how liberal an account of knowledge is, it will not allow for knowledge of not being a BIV. Hence, according to Dretske, there is only one successful way of rebutting this argument: to give up closure, that is, to claim knowledge of one's hands while admitting one can't know one is not a BIV. I disagree with that, and I will now proceed to explain why.

3. Why We Might Think We Don't Know We Are Not BIVs

Many epistemologists think that the third premise of the BIV Closure Argument is extremely plausible. The more one thinks about it, the more plausible it becomes. These epistemologists seem to be convinced that at least one of the following two theses is true:

1. 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.

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

One has some evidence for thinking one is *not* a BIV, but that evidence is not 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 used to have before envatment. Consequently, one cannot have any evidence for thinking one is not a BIV.⁶

Those who favor the Not-Good-Enough Thesis assert that we have *some* 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. Dretske would say that such evidential items do not add up to conclusive reasons. Therefore, whatever anti-BIV evidence we might have, it isn't good enough for us to know that the BIV hypothesis is false.

4. Easter Bunny Deception

Both theses strike me as false. For an assessment of their plausibility, I recommend considering a different type of skeptical scenario, one in which one is envatted or otherwise deceived not 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 agent of deception. Being handless, and indeed being altogether disembodied, the Easter Bunny's victim labors under the illusion of having a normal body. The Easter Bunny deception analog to the BIV argument, then, runs as follows:

⁶ This seems to be Dretske's view. See *Debates*, p. 23.

The Easter Bunny Deception Argument

- (1) $[KH \ \& \ K(H \rightarrow \sim EBD)] \rightarrow K\sim EBD$ ⁷
- (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 and hiding eggs. 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, then we don't think that knowledge requires truth-entailing evidence – which is to say we are endorsing fallibilism. According to fallibilist evidentialism, 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 nonexistence. 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 Easter Bunny Nonexistence Thesis

We have evidence for Easter Bunny 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 knowledge can be enlarged using deduction, then you should reject the No-Evidence-At-All and the Not-Good-Enough Thesis. 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

⁷ $K\sim EBD$ =I know there is no Easter Bunny deception (I know that I'm not deceived by the Easter Bunny).

evidence using deduction. You reason as follows:

The Easter Bunny Nondeception Closure Step

The Easter Bunny does not exist.

If so, I'm not deceived by the Easter Bunny.

Therefore:

I'm not deceived by the Easter Bunny.

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.

The We-Have-Evidence Thesis strikes me as true. It seems to me, therefore, that the No-Evidence-At-All Thesis is false.

Next, consider the Not-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 nonexistence, 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 Easter Bunny 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 Not-Good-Enough Thesis is false. Since the We-Know-It Thesis is exceedingly plausible, it seems to me it supplies us with a good reason to reject the Not-Good-Enough Thesis.

The arguments for the We-Have-Evidence and the We-Know-It theses presuppose that we

can expand our stock of knowledge using deduction. So they rely on *closure*. Dretske doesn't hold that we can never use deduction to expand our stock of knowledge.⁸ However, he would deny that, using Easter Bunny nonexistence as our starting point, we can, using deduction, come to know that we are not victims of Easter Bunny deception. Deduction does not generate knowledge when it comes to consequences for which we lack conclusive reasons. That's why, on Dretske's view, closure fails. What's nice about the deceptive Easter Bunny and other deceivers of the same ilk is the following: They illustrate just how painful it is to accept closure failure, at least when it comes to obviously non-existing agents of deception. For the position Dretske advocates implies conjunctions as abominable as 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) $K\sim EBE^9$
- (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 sum up: The Easter Bunny nonexistence 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.

⁸ See his essay in the Contemporary Debates volume.

⁹ $K\sim EBE$ =I know the Easter Bunny does not exist.

5. Rebutting the BIV Argument

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) $K\sim BIVE^{10}$
- (2) $K(\sim BIVE \rightarrow \sim BIV)$
- (3) $[K\sim BIVE \ \& \ K(\sim BIVE \rightarrow \sim BIV)] \rightarrow K\sim BIV$
- (4) $K\sim BIV$

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 nonexistence 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 nonexistence of BIVs. But surely it eliminates all reasonable doubt about their nonexistence. So, arguably, the proposition that there are no BIVs meets the standard of knowledge that fallibilist evidentialism endorses. We know that BIVs don't exist, and hence, performing an easy deduction, we can know that we are not victims of envatment.

Of course this could change. Perhaps at some point in the future, envatment will be a widespread and much reported phenomenon. Then it's going to be more difficult to know one is not a

¹⁰ $K\sim BIVE$ =I know that BIVs don't exist.

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 review the salient points. I think the best response to BIV skepticism goes like this. We have excellent evidence for thinking that BIVs don't exist, evidence that eliminates all reasonable doubt about BIV nonexistence. We have such evidence even though our overall evidence would be the same if we were BIVs. Therefore, the No-Evidence-At-All Thesis is false. Using BIV nonexistence as a premise, an easy deduction allows us to infer that, if BIVs don't exist, we are not BIVs. Anti-BIV-existence evidence thus becomes anti-BIV-deception evidence. But is our anti-BIV-deception evidence good enough to *know* we are not BIVs? Since the deduction in question is simple and rationally compelling, there is no reason to deny that we know we are not BIVs by deducing it from the nonexistence of BIVs. We should, therefore, reject the Not-Good-enough Thesis. If Dretske's conclusive reasons requirement were true, the inference from BIV nonexistence to one's not being a BIV would be blocked. However, Dretske's conclusive reasons requirement is not obviously true. To the contrary, it is highly problematic. If we reject it, no obstacle remains to saying that we can know we are not BIVs by deducing it from the nonexistence of BIVs.

6. Does Evidentialist Anti-Skepticism Beg the Question?

If there is one primary criticism my evidentialist response to skepticism invites, it is that I beg the question against skepticism. In this section, I will discuss whether this criticism sticks. Later, I'll examine a response that I think would be question begging.

My response to the BIV Argument appeals to the premise that I know BIVs don't exist. It might be thought that reliance on this premise begs the question against the skeptic, for the skeptic would of course argue that, since I don't know I'm not a BIV, I don't know that BIVs don't exist. So I think I must admit that I'm helping myself to a premise the truth of which the skeptic is not prepared to grant. It must not be overlooked, however, that the skeptic is using a premise the truth of which *I* am not prepared to grant, namely the premise that I don't know I'm not a BIV. It is of course right that, if I really don't know that I'm not a BIV, then I don't know that BIVs don't exist. On the other hand, it is also true that, if I really *know* BIVs don't exist,

then I do know I'm not a BIV. This symmetrical structure reveals itself when we compare the two arguments in question:

BIV Counter-Argument

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

Skeptical Reply

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

According to the anti-skeptical argument on the left, I know I'm not a BIV because I know BIVs don't exist. According to the skeptical reply on the right, I don't know BIVs don't exist because I don't know I'm not a BIV. The second premise of each of these arguments denies the conclusion of the other argument. What we are seeing is a clash between the skeptic's *modus tollens* and the anti-skeptic's *modus ponens*. So if there is any question-begging going on, it's mutual.

Consequently, when we ask which of the arguments is better, we need to look past the issue of question-begging and assess the plausibility of the premises. Since both arguments share the same first premise, we need to consider each argument's second premise. So the question comes down to this: What is more plausible, the skeptic's premise that I don't know I'm not a BIV, or the anti-skeptical premise that I know BIVs don't exist?

To answer this question, I suggest to consider two more pairs of arguments, thus generating an anti-skeptical slippery slope. Consider first how the skeptic would respond to the Easter Bunny Deception Counter-Argument:

The EB Deception Counter-Argument

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

Skeptical Reply

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

According to the argument against Easter Bunny deception, I know I'm not deceived by the Easter Bunny because I know the Easter Bunny does not exist. According to the skeptical response, I don't know that the Easter Bunny does not exist because I don't know that I'm not deceived by the Easter Bunny. Which argument is better? It seems to me when it comes to Easter Bunny deception, the skeptic's response doesn't look very good. The skeptical challenge seems silly, without any bite. We can let this decline continue by considering another, this time completely absurd skeptical argument. It goes like this: "If I know I have hands, I know I'm not deceived by a round square. But I don't know that I'm not deceived by a round square. Therefore, I don't know that I have hands." Now consider the following two arguments:

The RS Deception Counter Argument

- (1) $K \sim RSE \rightarrow K \sim RSD$
 (2) $K \sim RSE$
 (3) $K \sim RSD$

Skeptical Reply

- (1) $K \sim RSE \rightarrow K \sim RSD$
 (2) $\sim K \sim RSD$
 (3) $\sim K \sim RSE$

I say I know I'm not deceived by a round square because I know round squares don't exist, and the skeptic replies that I don't know round squares don't exist because I don't know I'm not deceived by a round square.

It is difficult to take Round Square Skepticism seriously. When we consider BIV skepticism, it is perhaps not easy to see whether the skeptic or the anti-skeptic has the upper hand. To clear this up, we can employ the device of a slippery slope. At the end of the slippery slope, the light of reason shines brightly. Here, we are observing a skeptical argument that can only be classified as preposterous. Round Square Skepticism is not a serious challenge. It is a mere play with words, utterly lacking in substance. Surely, the nonexistence of round squares is not something that can seriously be called into question. It seems to me, therefore, it would be rather odd to accuse the evidentialist response to Round Square Skepticism of begging the question.

Now consider BIV Skepticism, located at the other end of the slope. If BIV Skepticism is a serious challenge and evidentialist anti-skepticism begs the question against it, then something must have changed on the way up during the transition from Round-Square Skepticism to Easter Bunny Skepticism to BIV Skepticism. What might have changed? Perhaps it is the strength of the evidence. So perhaps the following diagnosis is true:

- 1 The evidence against the existence of round squares is as strong as it can get. There is no reasonable doubt as to the nonexistence of round squares. Therefore, round-square skepticism is unreasonable. The anti-skeptic does not beg the question against it.
- 2 The evidence against the Easter Bunny's existence is not quite as strong, but still very strong, eliminating all reasonable doubt. Therefore, Easter Bunny skepticism is unreasonable. The anti-skeptic does not beg the question against it.
- 3 But when it comes to evidence against the existence of BIVs, we witness a catastrophic weakening. The nonexistence of BIVs is very much in doubt. BIV skepticism is, therefore, reasonable, and the anti-skeptic begs the question against it.

Well, while I find myself in agreement with the first two steps of this diagnosis, I must reject the third. I don't think there is a catastrophic weakening of the evidence when it comes to the nonexistence of BIVs as opposed to the nonexistence of round squares or the nonexistence of the

Easter Bunny. Rather, I think there is still plenty of evidence making the anti-skeptic's position reasonable. So my view is that, while the strength of the evidence diminishes with each step on the slippery slope, it does so only in very small increments. No changes occur that dramatically change the picture when we move from Round Square Skepticism to BIV Skepticism. Hence I suggest to replace the third part of the diagnosis to:

- 3* When it comes to evidence against the existence of BIVs, no significant weakening has taken place. There is, therefore, no reasonable doubt as to the nonexistence of BIVs. BIV skepticism is, therefore, unreasonable, and the anti-skeptic does not beg the question against it.

Critics of evidentialist anti-skepticism who wish to oppose my line of reasoning are confronted with a dilemma. Consider the following thesis:

The No-Dramatic-Change Thesis

The slippery slope from Round Square to BIV Skepticism does not involve a dramatic weakening of the evidence.

Here is the dilemma critics of my slippery slope argument face: they must either accept or reject the No-Dramatic-Change Thesis. If they accept the thesis, it remains unclear why skepticism fails when it employs the deceiving round square hypothesis but is elevated to a serious challenge when it employs the BIV hypothesis. If they reject the thesis, they incur the obligation to identify a change significant enough to rehabilitate BIV skepticism given the obvious failure of Round Square skepticism. It is not easy to see, however, exactly where that change might occur.

7. Two Objections

It could be objected that a dramatic change does occur. It occurs at the transition from Round Square Skepticism to Easter Bunny Skepticism. Skeptical hypotheses, if effective, must be logically consistent, representing genuine possibilities. Round Square deception does not meet this requirement, and thus is not suitable for the skeptic's purposes.

There are two problems with this objection. First, if we say that a dramatic change occurs at the transition from Round Square to Easter Bunny Skepticism, we assign, as a consequence, too much credit to the latter. It is of course true that, unlike a deceiving round square, a deceiving Easter Bunny is at least logically possible. But that doesn't mean the Easter Bunny deception

hypothesis needs to be taken seriously. Though Round Square deception is even less plausible than Easter Bunny deception, certainly the difference is way too small to have the result of making Easter Bunny Skepticism look good by way of comparison.

Second, I think it would be a mistake to think a skeptical hypothesis is discredited simply because, if false, it is necessarily false. We are, after all, liable to make mistakes about what's logically possible and what's not. It strikes me as arbitrary, therefore, to say that a dramatic change takes place at the transition from Round Square to Easter Bunny Skepticism. *In principle*, it is possible for the skeptic to cast doubt on the reliability of our modal intuitions. Whether such doubts are to be taken seriously depends on the error possibilities the skeptic appeals to. If the skeptic were to present us with data suggesting that our modal intuitions are unreliable, then we *would* be confronted with a serious worry. If, however, the skeptic presents us with the hypothesis that we are deceived by a round square into thinking that round squares don't exist, then there is no reason to worry. But the explanation of why there is no reason to worry is not the shift from empirical to logical plausibility but rather the use of an utterly implausible error possibility. If a skeptical argument is to succeed, the skeptical hypothesis it is based on must itself be plausible or at least not be implausible. That is what matters. What does not matter is whether the skeptical hypothesis targets our empirical evidence or our modal intuitions.

According to the second objection, I am too confident in my negative assessment of Easter Bunny and Round Square Skepticism. Why not let the direction of the argument go the other way? Why not, that is, argue as follows: "Since I don't know I'm not a BIV, I don't know that BIVs don't exist. Since that's an unassailable starting point, epistemological reflection tells us we must employ analogous reasoning when it comes to a deceiving Easter Bunny or a deceiving round square. Starting with the plausibility of BIV skepticism, I should conclude that, since I don't know that I'm not deceived by the Easter Bunny, I don't know that the Easter Bunny does not exist, *and* that, since I don't know I'm not deceived by a round square, I don't know that round squares don't exist. So instead of undermining BIV Skepticism, what the slippery slope does is illustrate how strong and far-reaching the skeptic's position really is."

Here is my response to this objection. BIV skepticism rests on the premise that I don't know I'm not a BIV. This premise shouldn't dogmatically be turned into an article of faith. It's a problematic premise in need of support. If what can be said in support of it were beyond any reasonable doubt, then the objection would have a point. But what epistemologists have said in

support of this premise is debatable. *If* knowledge requires infallible evidence, or *if* knowledge requires a conclusive reason, or *if* knowledge requires sensitivity, then it follows of course that one doesn't know one is not a BIV. But none of these requirements are uncontroversial. Likewise, if the No-Evidence-At-All Thesis or the Not-Good-Enough Thesis were true, then, from an evidentialist point of view, it would be impossible to know that one is not a BIV. I have argued, however, that both of these theses are false. According to the evidentialist approach I have been advocating, knowledge results from the possession of evidence that eliminates all reasonable doubt. From that point of view, asserting that one is unable to know one is not a BIV is not an innocuous starting point. To the contrary, it is a deeply problematic claim that, all things considered, should be rejected.

8. What Would be Begging the Question?

I have argued that, when one derives knowledge of not being a BIV from knowledge of the nonexistence of BIVs, one does not beg the question against BIV skepticism. In this section, I will consider an argument that I am inclined to think does beg the question. Before I introduce it, let's look at an example that nicely illustrates the difference between begging and not begging the question. Suppose I believe about a table before me, on the basis of its looking red to me:

- (1) The table is red.

Suppose further someone confronts me with the following worry:

- (2) The table is white and illuminated by red light.

Suppose further I respond to this challenge by saying I know the table isn't white and illuminated by red light because I know it is red. That, I think, would be begging the question.¹¹ Knowledge of the table's redness is the very thing for which the red light hypothesis is a defeater. That's why appealing to knowledge of the table's redness is not the right response to defeating the red light hypothesis. In general terms, one cannot defeat a defeater by appealing to the very knowledge the defeater threatens. So if I wish to avoid begging the question, what I need to do is defeat the defeater by stating evidence in support of a defeater-defeater such as

- (3) There are no red lights present.

If I don't have evidence in support of (3), then the table's looking red does not put me in a

¹¹ I'm in agreement with Stewart Cohen on this. See his "Basic Knowledge and the Problem of Easy Knowledge," *Philosophy and Phenomenological Research* 65 (2002): 309-329.

position to know it is red. But if, on the basis of such evidence, I come to know that (3) is true, then the red light hypothesis is defeated, and the table's looking red gives me knowledge of its being red.

Now let's return to the question of how to respond to the BIV Argument. What is at issue is whether I know

(4) I have hands

on the basis of hand-like experiences. The BIV hypothesis,

(5) I am a BIV

is a defeater threatening my knowledge of (4). Suppose I respond to the threat (5) poses by saying I know I'm not a BIV because I know that I have hands. That, I think, would be question-begging, just like the red table response to the red light hypothesis would be question-begging. The I-know-I-have-hands response would be question-begging because one can't defeat a defeater by using a knowledge claim the legitimacy of which the defeater threatens. Consider my evidence for thinking I have hands: I can see and feel them. That I would have such evidence is of course just what's to be expected if I actually have hands. However, it's also just what's to be expected if I were a BIV. Until further bits of evidence come into play, the BIV hypothesis explains my hand evidence just as well as the fact of my having hands. Likewise, I can't defeat the red light hypothesis by saying I know the table is red. For if I have nothing but the table's looking red to go on, its being red is just as well explained by red light illumination as by its actually being red.

Now suppose I respond to the BIV hypothesis by claiming that, on the basis of the evidence cited earlier, I know

(6) BIVs don't exist.

Why would it be that claiming knowledge of (6) is not question-begging but an effective defeater for (5), whereas claiming knowledge of (4) begs the question and fails to be a defeater for (5)? I'm afraid there is no easy answer to that question, but here is what I think might be a satisfactory explanation. It is an essential ingredient of the BIV hypothesis that the envatted person takes herself *not* to be envatted. So a BIV labors under the illusion of having a normal body. Hence, as mentioned a moment ago, having hand-like experiences, falsely indicating knowledge of one's hands, is just what's to be expected if one were a BIV. That's why appeal to hand knowledge does not defeat the BIV hypothesis. In contrast, having plentiful evidence for

thinking that BIVs don't exist is *not* what's to be expected if one were a BIV. Whoever has been in charge of envatting the victim might as well have decided to let the BIV think that envatment is a common and wide-spread practice. So while having hand-like experiences is an essential element of the BIV hypothesis, having anti-BIV existence evidence is not. Arguably, that's why the BIV nonexistence response does not beg the question but is an effective defeater for the BIV hypothesis.

It could be objected that we can fine-tune the BIV hypothesis so as to make evidence for BIV nonexistence an essential ingredient of the revised hypothesis. And if BIV nonexistence evidence is an essential ingredient, it's just what's to be expected that as a BIV one would have such evidence. So consider the following skeptical alternative:

The Fine-Tuned BIV Hypothesis

I'm a BIV whose experiences, memories, and background beliefs are manipulated and fine-tuned so as to make me think there are good reasons for believing that BIVs don't exist.

According to the objection, the BIV nonexistence response begs the question against the fine-tuned BIV hypothesis, for that I would have BIV nonexistence evidence is now just what's to be expected if the hypothesis were true.

It seems to me, however, that we are now witnessing the limits of designing a non-defeasible skeptical hypothesis: a hypothesis against which there could not be a non-question-begging response. Consider the recipe behind the fine-tuned BIV hypothesis, which, I suggest, results in producing a self-defeating hypothesis because it comes with a built-in reason for rejecting the hypothesis:

The Self-Defeating Skeptical Hypothesis

A powerful agent of deception has seen to it that hypothesis H is true while giving me plenty of evidence for thinking that H is false.

According to the objection we are now considering, it would be question-begging to argue against hypothesis H on the basis of the anti-H evidence the deception master was friendly enough to supply me with. But I just don't think the argument would be question-begging. Rather, I think if the agent of deception wants to succeed in preventing me from having a good argument against thinking I'm deceived, he had better design an illusion without built-in evidence for thinking that I'm not subjected to any such illusion. If the deceiver prefers to supply me with such evidence, then he deceives me in such a way that it's reasonable of me to think

that I'm not deceived.¹²

There is, then, something special about BIV-nonexistence evidence. It is, by its very nature, the basis for a non-question-begging argument against the BIV hypothesis. Let's distinguish between BIVs who have such evidence, and BIVs who do not. BIVs who lack such evidence have no argument available for concluding that they are not BIVs. BIVs who have such evidence are justified in thinking that they are not BIVs. Likewise, we may distinguish between normal, non-envatted people who have BIV nonexistence evidence and those who do not. Unlike the latter, the former are in a position to give a non-question-begging argument for thinking that they are not BIVs. On the basis of such an argument, they can *know* that they are not BIVs.

¹² Which victim of radical deception is more strongly deceived: someone who has good reasons to think he is not deceived, or someone who has no clue as to whether he is deceived? Arguably, being deceived while being justified in thinking one is not deceived is the worse kind of being deceived.