

Matthias Steup

Easter Bunny Deception and the Failure of Skepticism

Let's begin with the familiar BIV argument:

- (1) If I don't know I'm not a BIV, then I don't know I have hands.
- (2) I don't know I'm not a BIV.
- (3) I don't know I have hands.

Philosophers who consider this argument sound are Skeptics, as opposed to Anti-Skeptics, who think that at least one of the premises is false. The kind of anti-skepticism I'm interested here denies the second premise.

Skeptics, obviously, take skepticism seriously. As a result, they choose to ask questions that, at least in a philosophical context, sound rather serious. For example, they might ask, with a certain inflection of voice intended to emphasize the seriousness of the question:

How do you know that you were not kidnapped by a Mad Scientist who turned you into a BIV having mental states and experiences exactly like those you are having right now?

Interestingly, though, they do not put the question like this:

How do you know that you were not kidnapped by the Easter Bunny who turned you into a BIV having mental states and experiences exactly like those you are having right now?

The second of these questions would not serve the skeptics as well as the first. For one reason, it would be hard to make the second question sound serious. Moreover, whereas the first question is likely to leave many people puzzled and confused, it's hard to imagine that confrontation with the second question will result in a fit of epistemological anxiety. Most people, I think, would not take that challenge seriously at all. They would without hesitation claim to know that the Easter Bunny does not exist, and hence would claim to know that they were not abducted and envatted by the Easter Bunny.

The Easter Bunny version of the BIV hypothesis is one the skeptics have never considered. But now that it's on the table, we might as well compare it with the Mad Scientist version and examine whether choosing the Mad Scientist version safeguards skepticism from the failure that threatens when the envatting agent is not a Mad Scientist but the Easter Bunny. Let's describe the two hypotheses as follows:

The Mad-Scientist Hypothesis

At some point in the past I was abducted by a Mad Scientist who envatted me and now stimulates my brain in such a way that a perfect illusion of an ordinary life is generated. The experiences and mental states I have in this possible scenario are identical to those I have in the actual world.

The Easter-Bunny Hypothesis

At some point in the past I was abducted by the Easter Bunny who envatted me and now stimulates my brain in such a way that a perfect illusion of an ordinary life is generated. The experiences and mental states I have in this possible scenario are identical to those I have in the actual world.

The two hypotheses are identical except for the Mad Scientist/Easter Bunny difference. The question is whether there is any significant *epistemic* difference between them. Next, I will set forth an anti-skeptical argument based on the premise that the two hypotheses are of comparable epistemic status. For that purpose, let's consider two versions of the BIV Argument:

The Mad Scientist Argument

- (1) If I don't know that the Mad Scientist Hypothesis is false, then I don't know I have hands.
- (2) I don't know that the Mad Scientist Hypothesis is false.
- (3) I don't know I have hands.

The Easter Bunny Argument

- (1) If I don't know that the Easter Bunny Hypothesis is false, then I don't know I have hands.
- (2) I don't know that the Easter Bunny Hypothesis is false.
- (3) I don't know I have hands.

Using the Easter Bunny Argument as a hostile analogy, we can now argue against the Mad Scientist Argument as follows:

The Analogy Argument

- (1) The Easter Bunny Argument is unsound because its second premise is false.
- (2) The second premise of the Easter Bunny Argument and the second premise of the Mad Scientist Argument are epistemically alike.
- (3) The Mad Scientist Argument is unsound because its second premise is false.

In response to the Analogy Argument, the Skeptics face two options:

(i) Denial of the first premise, acceptance of the second.

(ii) Acceptance of the first premise, denial of the second.¹

Option (i) meets with the following difficulty: Knowing that the Easter Bunny does not exist, I can use closure to deduce, and thus come to know, that the Easter Bunny Hypothesis is false. To block that move, the Skeptics would have to deny either closure, or fallibilism, or that the Easter Bunny's non-existence is something I know. None of these seem attractive.

Option (ii) raises a serious problem as well. It's not easy to see why there should be a significant epistemic difference between the Easter Bunny's non-existence and the non-existence of Mad Scientists. Once it's agreed that we know that the Easter Bunny does not exist, it's not exactly an enormous leap when we then also claim to know that Mad Scientists, at least of the abducting and envatting kind, do not exist. In what follows, I will examine each option more closely.

The First Premise

Suppose, then, the Skeptics reject the first premise of the Analogy Argument. If so, the Skeptics claim that

I don't know I wasn't envatted by the Easter Bunny is true. As I mentioned a moment ago, we might now wonder whether the Skeptics also want to assert:

I don't know that the Easter Bunny does not exist.

Presumably not. But then we might reason as follows. First, let's stipulate that 'EB' means 'The Easter Bunny exists', and 'EBH' refers to the Easter Bunny Hypothesis, according to which the Easter Bunny abducted and envatted me. Second, we note that EBH entails EB. Third, we appeal to closure:

Modus Tollens Closure

$[K\sim q \ \& \ K(p \rightarrow q) \ \& \ CD^2] \rightarrow K\sim p$.

Applying this to the issue at hand, we get

Closure of Easter Bunny Deception

$[K\sim EB \ \& \ K(EBH \rightarrow EB) \ \& \ CD] \rightarrow K(\sim EBH)$.

If the Skeptics accept Modus Tollens Closure, they are not in a position to claim that, while I know that the Easter Bunny does not exist, I do not know

that I was not abducted and envatted by the Easter Bunny. Closure opponents, of course, will make precisely that claim. But that means they get stuck with one more abominable conjunction. The abominable conjunction DeRose identified is this: I know I have hands, but I don't know I'm not a handless BIV. It seems to me the conjunction "I know the Easter Bunny doesn't exist, but I do *not* know I wasn't abducted and envatted by the Easter Bunny" is no less abominable.

Let us back up one step. Dretske and Nozick reject closure because they think that, if the skeptic is permitted to use closure as a premise, the skeptic wins. It seems to me, however, that actually the closure principle helps not the skeptic but the anti-skeptic. If closure holds, the anti-skeptic can reason from an obvious and uncontroversial falsehood that a skeptical hypothesis entails to the falsehood of that skeptical hypothesis. If that's right, then the aim of escaping skepticism is not a good reason for closure denial, especially considering the enormous costs associated with that move.³

Let's return to those Skeptics who accept closure. Since they accept closure, they can reject the first premise of the Analogy Argument only if they also reject the claim that we know that the Easter Bunny does not exist. Some Skeptics might want to do just that. So do we really know that the Easter Bunny does not exist?

Fallibilism and Infallibilism

I think we know the Easter Bunny does not exist. Obviously, though, my evidence for the Easter Bunny's non-existence does not *entail* the Easter Bunny's non-existence. So if we work with a Cartesian concept of knowledge that requires infallibility, we can't be said to know that the Easter Bunny doesn't exist. But, when I say we know the Easter Bunny does not exist, what I have in mind is not knowledge of the Cartesian kind in mind but rather fallibilist knowledge. Let us define this notion as follows:

Fallible Knowledge

Knowledge of p requires evidence for p, but the required evidence for p need not entail p.

¹ A third option would be to deny both premises.

² CD = competent deduction.

³ See Hawthorne. Of course there might be a prior theoretical commitment that's not negotiable in the present context, but independently motivated. Dretske: K requires conclusive reasons. We have a conclusive reason for the EB, but not a conclusive reason for EEB. So Dretske is unlikely to endorse closure.

Fallibilists might disagree on how strong a person's evidence for p must be to be strong enough for knowledge of p. According to the kind of fallibilism I'm endorsing here, knowledge of p requires evidence that elevates p to the lofty but not excessively lofty status of being beyond a reasonable doubt.

So if we know the Easter Bunny does not exist, we must have evidence that eliminates any reasonable doubt as to the Easter Bunny's non-existence. What is that evidence? I will mention only this much: distributing and hiding Easter eggs in all the countries in which that custom exists, either single-handedly or assisted by an army of helpers, requires intellectual and organizational skill bunnies don't have. That reason, by itself, suffices for knowing that the Easter Bunny does not exist. Additional evidence is easy to find.⁴ The total body of Easter-Bunny-non-existence evidence is of such overwhelming strength that, it seems to me, the Skeptics are well-advised if they agree that, given a fallibilist understanding of knowledge, we know that the Easter Bunny does not exist.

The Skeptics could argue, however, that a fallibilist concept of knowledge makes no sense. Knowledge, by definition, requires infallible evidence, or, to put the point in more general terms, the exclusion of all error possibilities. That's the view David Lewis holds. Here's what he says on this matter:

It seems as if knowledge must be by definition infallible . . . To speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just sounds contradictory. (549)

And:

We are caught between the rock of fallibilism and the whirlpool of skepticism. Both are mad! Yet fallibilism is the less intrusive madness . . . We can get used to it, and some of us have. No joy there -- we know that people can get used to the most crazy philosophical sayings imaginable. If you are a contented fallibilist, I implore you to be honest, be naive, hear it afresh. 'He knows, yet he has not eliminated all possibilities of error.' Even if you've numbed your ears, doesn't this overt, explicit fallibilism *still* sound wrong? (550)

I take it the argument suggested in these passages is the following: "If we are fallibilists, we will utter knowledge ascriptions that sound wrong.

⁴ Paws. None were ever observed. Location? Language?

Therefore, fallibilism must be rejected as a conceptual confusion." Next, I'll respond to this argument. First, though, I should mention that Lewis rejects evidentialism and instead works with an externalist error-possibility-elimination approach. It seems to me, though, that Lewis's basic point remains the same irrespective of which approach is chosen. I will, therefore, treat Lewis's demand for infallible elimination of error possibilities as the demand for infallible evidence.

Let's look, then, at some of the sentences we might assert if we are fallibilists:

S1 I know that that the Easter Bunny does not exist, but it is logically possible that the Easter Bunny exists.

This doesn't sound wrong to me. In general terms, S1 merely makes the point that what we know is not restricted to necessary truths.

S2 I know that the Easter Bunny does not exist on the basis of evidence E, but there are possible worlds in which I am in possession of evidence E and the Easter Bunny exists.

Perhaps I don't hear this afresh enough, but S2 doesn't sound mistaken to me at all. It rather sounds to me like just the sort of thing we should say about our knowledge of the Easter Bunny's non-existence. The third sentence tightly follows the wording of Lewis's example.

S3 I know that the Easter Bunny does not exist, but I have not eliminated all possibilities of error.

'Eliminating error possibilities' is a locution in need of clarification. Let's conceive of it in evidentialist terms. This gives us two options:

S4 I know that the Easter Bunny does not exist, but my evidence does not entail the Easter Bunny's non-existence.

S5 I know that the Easter Bunny does not exist, but my evidence does not eliminate reasonable doubt about the Easter Bunny's non-existence.

S5 does sound wrong to me. But S5 is not entailed by fallibilism as I understand it here. S4, on the other hand, is just what a fallibilist should say. But, once again, even while sincerely attempting to be naive and hear S4 afresh, it doesn't sound like a mistake to me.

The Second Premise

Let's move on to the second premise. The question we need to address is this: Is our evidence for the non-existence of envatting Mad Scientists

analogous to our evidence for the non-existence of the Easter Bunny, in the following respect: Does it eliminate any reasonable doubt as to the non-existence of Mad Scientists? I'm inclined to say it does. There is what we might call 'common knowledge.' It's the sort of stuff one initially learns going to school and later learns by reading books, newspapers, magazines, or surfing the internet. Here are two things I think I know on the basis of such common knowledge. First, I know that there isn't a nuclear bomb in my basement. Second, I know that there isn't a million dollars hidden in my mattress. In either case, my knowledge is inferential. It's not based on direct perception. I have perceptual access right now to neither my basement nor the inside of my mattress. Nevertheless, I know the two things I just mentioned. I know them because our common knowledge addresses the location and treatment of both large sums of money and nuclear explosives. In each case, what I know about these issues removes any reasonable doubt. Envatting Mad Scientists, I suggest, fall into the same category, just as creatures like Big Foot, the Abominable Snowman, and the Loch Ness Monster. You know that Mad Scientists of the envatting kind do not exist in the same way in which you know that the creatures I just mentioned do not exist, or that there isn't a million dollars in your mattress, or that there is no nuclear bomb in your basement.

So let's say you know that envatting Mad Scientists do not exist. Obviously, the Mad Scientist Hypothesis entails the existence of at least one. Thus, performing a competent closure-type deduction, you can come to know that the Mad Scientist Hypothesis is false. And that's just the point of the Analogy Argument: the Mad Scientist Argument fails because, just as we know that there are no Easter Bunnies, we know that there are no Mad Scientists.

I'm inclined to think, then, that the second premise of the analogy argument is true. We know that abducting and envatting Mad Scientists don't exist, just as we know that the Easter Bunny does not exist. Therefore, we know that the Mad Scientist Hypothesis is false.

Closure-Based Anti-Skepticism

It would be nice if the strategy of using closure against the skeptic were applicable to all skeptical scenarios. Now, there are of course many possible worlds in which that strategy won't work. For example, there are worlds in which abducting and envatting people is a wide-spread and popular

practice. In such worlds, the needed premise asserting the non-existence of envatting agents is not available. It seems to me, though, that the *actual* world is very different from possible worlds in which deception-inducing agents abound. The actual world, I would like to think, is such that, if we know that p, then every skeptical alternative to p does indeed entail the existence of something that we know, on the basis of evidence removing all reasonable doubt, not to exist or not to obtain.

Let's review a few more alternatives. Here is Lewis again:

Let your paranoid fantasies rip -- CIA plots, hallucinogens in the tap water, conspiracies to deceive, old Nick himself -- and soon you find uneliminated possibilities of error are everywhere. (549)

Given what Lewis means by eliminating error possibilities, it is indeed the case that we cannot eliminate such skeptical scenarios.⁵ We are now, however, looking at these possibilities from the point of view of evidentialist fallibilism. To eliminate these possibilities, what we need is evidence good enough to remove any reasonable doubt as to their non-actuality. Let's look at each in turn.

Going in reverse order, there is, to begin with, old Nick: the devil himself. Well, let's just say there's no harm in dismissing the devil with no further ado and move on.

Next, there are conspiracies to deceive. That's a bit general. Unless there are plausible details, taking a conspiracy seriously is not good epistemology but an indication of paranoia.

Third, hallucinogens in the tap water. Here, the idea is that, because of tap-water delivered hallucinogens, I am subjected to a massive, BIV-type illusion about my actual situation. Obviously, this Hallucinogens-in-the-Tap-Water Hypothesis entails the existence of hallucinogens in the tap water -- not just of any kind, but hallucinogens powerful enough to generate a BIV delusion of a normal life when reality is radically different from what it appears to be. I would say we know that such hallucinogens don't exist.

Finally, CIA plots. Now we are to consider the CIA Plot Hypothesis, according to which my evidence is what it is now, but, because I'm a guinea pig in some CIA lab, my belief that (say) I have hands is false. This hypothesis entails several relevant propositions:

⁵ Hence the contextualist move of rescuing ordinary knowledge by claiming that in ordinary contexts, such scenarios are irrelevant.

- D There are envatment involving CIA plots.
- E The technology for successful envatment exists.
- F The know-how for successful envatment exists.

I'll make this short. What I call *successful envatment* involves two crucial aspects: First, keeping a brain alive and functioning for an extended period; second, generating the perfect illusion of a normal life by stimulating nerve endings. It's common knowledge that envatment, thus understood, belongs to the realm of the non-existent. So we know that E and F are false, and therefore know that D is false.

Let's dwell on this point just a bit longer. Suppose you develop an interest in ants, and you want to know whether a certain kind of ant exists in the region in which you live. To find out, you call up a colleague of yours in the Biology Department who is an expert on ants. He tells you that the type of ant in question occurs only in Africa. Now you know something you didn't know before. Likewise, if you feel unsure about envatment, you can

consult experts on neurophysiology. They will tell you that neurophysiology isn't there yet. As of today, envatment is not physical possible. Nor will it become physically possible any time soon. So by consulting an expert on neurophysiology, one can relatively easily come to know that E and F are false. But if E and F are false, so is D. And if we know that D is false, then we know that the CIA Plot Hypothesis is false.

In general terms, the anti-skeptical strategy I'm advocating is this: Skeptical alternatives lacking details don't have bite. Detailed ones entail the existence of things we know not to exist or not to obtain. Using such knowledge of non-existence and performing closure-type deductions, we can, given what the actual world is like, come to know that skeptical alternatives are false.