

DO YOU KNOW THAT YOU ARE NOT A BRAIN IN A VAT?

Ned MARKOSIAN

ABSTRACT: The topic of this paper is the familiar problem of skepticism about the external world. How can you know that you are not a brain in a vat being fooled by alien scientists? And if you can't know that, how can you know anything about the external world? The paper assumes Evidentialism as a theory about justification, and then argues that you are justified in believing that you are not a brain in a vat, in virtue of the fact that your evidence supports that belief. The paper also considers a number of different objections to this proposal. The upshot is that you do know that you are not a brain in a vat, and that you also know lots of things about the external world.

KEYWORDS: skepticism, external world, brain in a vat, evidentialism

1. Introduction

The topic of this paper is the familiar problem of skepticism about the external world. How can you know that you are not, for example, a brain in a vat being fooled by alien scientists? And if you can't know that, then how can you know anything about the external world? My aim is to spell out this problem in its strongest, most annoying form, and then to show how it can be solved. The upshot, I will conclude, is that you *do* know that you are not a brain in a vat, and that you do know lots of things about the external world. But let's start with the skeptical problem.

Imagine a human brain in a vat, with tubes and wires hooked up to it. The tubes are for blood flow in and out of the brain, to keep it alive and healthy. The wires are for sensory input, and for monitoring the brain's attempts to move its own body. All of the wires are hooked up to a fancy computer, which ensures that the brain's experiences are completely realistic. Imagine that it seems for all the world to that brain that it is a normal brain in a body, in a world with other humans, and trees, birds, etc.

Now consider this skeptical hypothesis.

The Brain In a Vat Hypothesis (BIVH): I am just a brain in a vat being fooled by alien scientists.

It would of course be pretty silly to believe BIVH. After all, you have no special reason to think it is true.¹ But can you know that BIVH is false? Can you know that you are not a brain in a vat? The problem is that all of your experiences are consistent with your being a brain in a vat. Indeed, if you were a brain in a vat, you would be having experiences exactly like the experiences that you are in fact having. So even assuming that you are not a brain in a vat (an assumption that the skeptic is happy to make), it seems like you can't *know* that BIVH is *false*.

Now, unfortunately, pretty much all of your beliefs about the external world are based on the assumption that BIVH is false. (By 'the external world' I just mean the world outside of your own mind.) So, for example, you believe that your best friend is real. But if you are just a brain in a vat, then your best friend is not real. Thus, your belief that your best friend is real is based on the assumption that you are not just a brain in a vat. Similarly with your belief that you have hands, that your bicycle is where you left it, that Barack Obama is President of the US, that grass is green, etc. All of these beliefs are based on the assumption that you are not a brain in a vat. So if you can't know that BIVH is false, then it looks like you can't really know any of your beliefs about the external world.

Here is the argument suggested by these considerations.²

The Brain In a Vat Argument

- (1) I don't know that BIVH is false.
- (2) If (1), then I don't know anything about the external world.

- (3) I don't know anything about the external world.

The problem is that The Brain In a Vat Argument is valid, the premises both seem true, and yet we want to deny the conclusion. Something has to give.

¹ What would count as a reason to think BIVH is true? Perhaps experiences like these: You are walking down the street, hearing the ordinary soundtrack of life – birds chirping, the wind in the leaves, etc. Then all of a sudden there is a needle-on-vinyl-like glitch in the soundtrack, immediately followed by a strange, alien voice saying, "How is the experiment proceeding with the human brain in a vat?" This is followed by another strange, alien voice responding, "The experiment is proceeding very well... Hey, don't lean on that button. Get your tentacle off that button!" And this is abruptly followed by a return to the soundtrack of ordinary life.

² Hilary Putnam made this kind of argument famous. See Hilary Putnam, *Reason, Truth and History* (Cambridge: Cambridge University Press, 1981), Ch. 1. Putnam has his own strategy for dealing with the argument, which I will not discuss here. (But see Jane McIntyre, "Putnam's Brains," *Analysis* 44 (1984): 59-61, for an excellent summary and criticism of Putnam's strategy.) For yet another anti-skeptical strategy see Jessica Wilson, "The Regress Argument Against Cartesian Skepticism," *Analysis* 72 (2012): 668-673.

2. Knowledge and Justification

Before we get to my proposed solution, it will be useful to make explicit some background assumptions that are relevant to our problem. Philosophers generally agree that knowledge involves three things: *belief*, *truth*, and *justification*. Here is an analysis of the concept of knowledge that is often taken as a starting point in epistemology.³

The JTB Analysis of Knowledge: *S* knows *p* if and only if (i) *S* believes *p*, (ii) *p* is true, and (iii) *S* is justified in believing *p*.

It turns out that there is a well known problem with The JTB Analysis. Edmund Gettier showed in 1963 that belief, truth, and justification are not *sufficient* for knowledge.⁴ But The JTB Analysis will do for our purposes, since it is not a controversial matter that *belief*, *truth*, and *justification* are all *necessary* conditions for knowledge, and as we'll see very shortly, that is what is relevant to our skeptical problem. (In other words, we can replace 'if and only if' in The JTB Analysis with 'only if' and we will have something true that is relevant to our skeptical problem.)

Once we have seen that knowledge always involves a justified, true belief, we can see why it appears that I don't know that BIVH is false. It's not that I don't *believe* that the BIVH is false, for I do. And it's not that the BIVH is not *false*, for it is. (Or so, at any rate, we will assume. And remember, the skeptic is happy to grant this assumption. i.e., the fact that BIVH is actually false is not the issue. (Although I will come back to this point later.)) No, the problem is that I have a true belief that does not appear to be *justified*. And the reason my belief that I am *not* a brain in a vat does not appear to be justified is that things seem to me exactly the way they would seem if I *were* just a brain in a vat. I.e., all of my experiences are consistent with BIVH.

Upshot: The reason it seems like I don't know that BIVH is false is that my true belief that it is false appears to be unjustified.

3. An Objection to Premise (1) of The Brain in a Vat Argument

I reject premise (1) of the argument. The premise says I don't know that BIVH is false, so to reject the premise is to say that I *do* know that BIVH is false. Basically,

³ Something like The JTB Analysis is often attributed to Plato. See his *Meno* 98. For an alternative approach, see Carrie Ichikawa Jenkins, "Knowledge and Explanation," *Canadian Journal of Philosophy* 36 (2006): 137-164.

⁴ See Edmund Gettier, "Is Justified True Belief Knowledge?" *Analysis* 23 (1963): 121-123.

Ned Markosian

then, to reject premise (1) is to say that I know that I am a normal brain in a body. Let's give a name to that proposition.

N = the proposition that I am a normal brain in a body.

As noted above, I believe N. And the skeptic grants that N is true. So here is the situation so far with respect to N.

✓ I believe N.

✓ N is true.

? My belief that N is justified.

Do I know N? Given our assumptions, that depends on whether my belief that N is justified. So to attack premise (1) of The Brain In a Vat Argument, I must argue that my belief that N is justified. (And it is worth emphasizing here that the skeptic agrees with this assessment of the situation. That is, the skeptic agrees with putting those first two check marks there.)

Now, given that what is at issue between the skeptic and me is whether my belief that N is justified, we must consider the question, *What is justification?* I.e., *What conditions must be satisfied in order for it to be true that some person, S, is justified in believing some proposition, p?* Here are some of the main theories that have been proposed by epistemologists in answer to this question.

Foundationalism: *S*'s belief that *p* is justified if and only if either (i) *S*'s belief that *p* is indubitable or else (ii) *S* is able to derive *p* from some other, indubitable belief(s).⁵

Coherentism: *S*'s belief that *p* is justified if and only if *S* has a coherent system of beliefs and *S*'s belief that *p* fits nicely into that system.⁶

Reliabilism: *S*'s belief that *p* is justified if and only if that belief is produced by a reliable belief-forming process.⁷

Evidentialism: *S*'s belief that *p* is justified if and only if that belief is supported by *S*'s total evidence.⁸

⁵ See for example, René Descartes, *Meditations on First Philosophy*. I am using 'indubitable' as a placeholder for whatever characteristic the Foundationalist identifies as the one that makes foundational beliefs justified.

⁶ See Erik Olsson, "Coherentist Theories of Epistemic Justification," in *The Stanford Encyclopedia of Philosophy* (2012). <http://plato.stanford.edu/entries/justep-coherence/>.

⁷ See Alvin Goldman, "Reliabilism," in *The Stanford Encyclopedia of Philosophy* (2008). <http://plato.stanford.edu/entries/reliabilism/>.

⁸ See Earl Conee and Richard Feldman, "Evidentialism," *Philosophical Studies* 48 (1985): 15–34.

Do You Know That You Are Not a Brain In a Vat?

I vote for Evidentialism, although I won't argue for it here. But I will argue, assuming Evidentialism is true, that my belief that I am not a brain in a vat is justified. If you like, you can take what I say as a big conditional: If Evidentialism is true, then my belief that I am not a brain in a vat is justified. I think that what I will say would also work on the assumption that Coherentism is true, as well as on the assumption that Reliabilism is true, and I will say a little bit about that below. But I will focus mainly on the Evidentialist version of the proposal.

So what does it mean for a belief to be *supported* by your evidence? Like other Evidentialists, I offer no official definition, but here are some examples to help clarify the idea.

- You have read in the newspaper online that the Orioles won the game, and you have also seen reports on TV saying the same thing, and no one has told you otherwise. Your evidence *supports* the belief that the Orioles won the game.
- You feel a tremendous pain in your head. Your evidence *supports* the belief that you are in pain.
- Tim is on his way to ask Maria out again, for the hundredth time, even though she has made it abundantly clear that she has no interest in going out with him. But Tim is optimistic, and believes that this time she will say Yes. Tim's belief that Maria will say Yes is *not supported* by his evidence.
- You have known Tom for 15 years, and he has always been a good, reliable, honest friend. Now he is earnestly telling you that he has no idea what happened to the money you left on the table. So you believe him. But, alas, it just so happens that he is lying to you this time. Still, your evidence that Tom did not take the money *supports* your belief that he did not do it. (This is one of those cases in which your evidence is misleading. Life is hard.)
- You're on the jury in a big murder trial. The first witness is an extremely reliable member of the community, she has nothing to gain from the defendant's being convicted, she says that she saw the defendant commit the murder, and she has a very convincing story. Similarly for a dozen other self-proclaimed eye-witnesses. Meanwhile, the DA presents tons of physical evidence – DNA test results, fingerprints, video footage, etc. – linking the defendant to the crime. Finally, the defendant takes the stand and denies doing it, while telling some ridiculous story. In this case you have *some* evidence that the defendant is innocent (namely, his own testimony), but it is outweighed by the overwhelming evidence against the guy, so that your *total evidence* supports the belief that he did it.

So much for the crucial notion of support. What about evidence? What is that?

I think that each one of us has, at any given time, a great deal of evidence. But I think there are three main kinds. First, and perhaps mostly, there are sensations. These come in the form of visual sensations, auditory sensations, tangible sensations, etc. They are the sights, sounds, smells, tastes, and touches of ordinary (and extraordinary) life. Sometimes they are called appearances. It might also make sense to call them sensory seemings.⁹

Second, there are *a priori* intuitions. It just seems true to me that two plus two equals four, that killing innocent children for no reason is morally wrong, and that a spatially discontinuous object would have to have proper parts. *A priori* intuitions are seemings, too, although of a different kind from sensory seemings.

And third, there are memory impressions. When I introspect and try to examine my own memory impressions, it seems like they come in two main forms. Sometimes they are like vague and fleeting sensations. When I consider the question of where my bicycle is located, a sort of incomplete, flickering image comes into my mind: an image of my bicycle in my garage. But other times memory impressions are more like *a priori* intuitions about certain propositions being true. Right now my memory tells me that I moved to Bellingham in 1998, but it's not the case that I have some vague and fleeting memory impression placing that event in 1998. (I do have fleeting memory impressions of pulling up in a rental truck. But nothing about those impressions places the event in 1998.) Instead, it's more like the relevant proposition is just in the *Remembered Stuff* folder in my mind, and I can tell that by introspection.

Perhaps I have said enough about evidence and support for us to return to our main question: Is my belief that N (the proposition that I am a normal brain in a body) supported by my evidence? I say Yes, and the skeptic says No. Moreover, it is worth emphasizing that this is the very question about which the skeptic and I disagree (on the assumption that Evidentialism is the correct theory of justification, that is).

Here is why I think my belief that N is supported by my evidence. To begin with, I don't have any memory impressions that bear on this point. (But it is

⁹ For discussions of a number of issues having to do with the nature of evidence from our senses, see G.E.M. Anscombe, "The First Person," in her *Metaphysics and the Philosophy of Mind* (Minneapolis: University of Minnesota Press, 1981); and Susanna Schellenberg, "Experience and Evidence," *Mind* (forthcoming). On the more general topic of self-knowledge, see Brie Gertler, "Self-Knowledge," in *The Stanford Encyclopedia of Philosophy* (2008). <http://plato.stanford.edu/entries/self-knowledge/>.

worth noting that this is largely because I don't have any memory impressions of the kind of glitzy soundtrack involving alien voices mentioned in footnote 1 above.) Nor do I have any relevant *a priori* intuitions. (After all, the question of whether I am a brain in a vat or a brain in a body appears to be a strictly contingent matter, and not something that could be decided from the armchair.) So it comes down to sensations.

Now, let me just check the evidence from my senses here. (Think of this paragraph as the live-blogging portion of the paper.) As I hold my hands out and look at them carefully, my visual sensations forcefully suggest that I have hands and arms. I rub my hands together and pat my arms, and as I do, my tangible sensations suggest the same thing. Clapping my hands lightly results in audible sensations of having hands, and when I hold one hand in front of my nose I get olfactory sensations of having a hand. (Nice soap!) Finally, putting one knuckle gingerly into my mouth produces taste sensations of having a hand. (Soap flavored, but unmistakably a hand.) Meanwhile, I have had no evidence at all (sensory or otherwise) of alien-scientist glitches. So my evidence overwhelmingly supports the proposition that I have hands, and arms, and a body.¹⁰

Evidentialism says that a belief is justified if and only if it is supported by the agent's evidence. Assuming Evidentialism, then, we have this situation.

N = the proposition that I am a normal brain in a body.

✓ I believe N.

✓ N is true.

✓ My belief that N is justified.

I think it is obvious that my evidence supports the belief that I am a normal brain in a body, since that is what my evidence makes it seem like to me. (And I think the same is true of your evidence and your belief that you are a normal brain in a body.) But here is an additional argument for this claim. Consider two brains:

¹⁰ Similar moves have been made, or at least considered, by various other philosophers. See, for example, John Pollock, *Knowledge and Justification* (Princeton: Princeton University Press, 1974), Ch. 2-5; William Alston, *Epistemic Justification* (Ithaca: Cornell University Press, 1989), Essays 1-3; Robert Audi, *The Structure of Justification* (Cambridge: Cambridge University Press, 1993), Ch. 3, 4, 10, and 12; James Pryor, "The Skeptic and the Dogmatist," *Nous* 34 (2000): 517-549; Earl Conee and Richard Feldman, *Evidentialism* (Oxford: Oxford University Press, 2004), 300-301; Michael Huemer, *Skepticism and the Veil of Perception* (Lanham: Rowman & Littlefield, 2001), Ch. V; and Berit Brogaard, "Phenomenal Seemings and Sensible Dogmatism," in *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, ed. Chris Tucker (Oxford: Oxford University Press, forthcoming). For criticism of this general idea, see Roger White, "Problems for Dogmatism," *Philosophical Studies* 131 (2006): 525-557.

mine, which is a normal brain in a body, and Sally's, which is a brain in a vat. Now, the whole point of the brain-in-a-vat scenario is that it seems to the brain in the vat that it is a normal brain in a body. So it *looks* to Sally like she is a brain in a body, it *feels* to her like she is a brain in a body, etc. Sadly, Sally's evidence happens to be misleading. *But if it is misleading, then it is leading somewhere.* And where it is leading is toward the belief that she is a normal brain in a body. That's the belief that Sally's evidence supports. Thus Sally is justified in believing that she is a normal brain in a body.¹¹ And my evidence is just like Sally's in terms of where it is leading. My evidence is leading toward the belief that I am a brain in a body. So Sally and I both have justified beliefs (in her case, that she is a normal brain in a body; in my case, that I am a normal brain in a body).¹² In each case, the belief is justified by the agent's evidence. In my case, the belief is also true. So I have knowledge.¹³ In Sally's case, the evidence is misleading and her belief is false. So she does not have knowledge.

Here is the situation with respect to Sally's belief that she is a normal brain in a body.

N2 = the proposition that Sally is a normal brain in a body.

✓ Sally believes N2.

✗ N2 is true.

✓ Sally's belief that N2 is justified.

The upshot is this: Brains in vats don't know that they are brains in bodies (because they have justified but false beliefs in the relevant propositions), but brains in bodies do know that they are brains in bodies (because they have justified, true beliefs in the relevant propositions).

Now, I think similar remarks apply with respect to Coherentism. I have a coherent system of beliefs, and my belief that I am a normal brain in a body fits nicely into that system. So, according to Coherentism, my belief is justified. (Notice that according to Coherentism, Sally also has a justified (but in her case false) belief.)

¹¹ Notice that the only way to resist this argument is to say that Sally's evidence is not misleading. And that's crazy. Poor Sally is the poster child for misleading evidence!

¹² Huemer makes a similar point about a brain in a vat in *Skepticism and the Veil of Perception*, Ch. V, Section 3.

¹³ This is a little bit hasty, in light of the Gettier problem. But I am assuming that the skeptic is not resting her case on some tricky, Gettier-style consideration.

The idea seems to work with Reliabilism as well. My belief that I am a normal brain in a body is the result of a process involving my brain and my sensory organs, and this process is quite reliable. So, according to Reliabilism, my belief is again justified. (Notice that according to Reliabilism, however, Sally does not have a justified belief. (For Sally's belief-forming process is, unfortunately, highly unreliable.) So Reliabilism, unlike Evidentialism and Coherentism, entails that a brain in a vat does not have justified beliefs about the external world.)

It's time for a big-picture moment. According to what I have said, I know that I am a normal brain in a body, because my belief that N is supported by my evidence. And for similar reasons I know that BIVH is false. So premise (1) of The Brain In a Vat Argument is false. And so I know lots of things about the external world.¹⁴

Now let's turn to some likely objections.

4. Objections and Replies

Objection 1: "You say that your evidence supports your belief that you are a brain in a body. *But your evidence equally supports the proposition that you are a brain in a vat.* After all, your evidence is perfectly consistent with BIVH!"

¹⁴ I want to emphasize that my response to the skeptical challenge posed by The Brain In a Vat Argument is quite different from the famous Moorean response to such challenges. For the Moorean response (as applied to this version of the skeptical challenge) is to accept the conditional 'If I don't know that I am not a brain in a vat, then I don't know that I have hands,' and to *tollens* where the skeptic wants to *ponens*. That is, the Moorean wants me simply to accept as a premise the proposition that I know that I have hands. And the Moorean typically claims that I do not need any justification for the relevant claim because it is more plausible than any premise in any skeptical argument. My approach, on the other hand, is to accept that beliefs like my belief that I have hands are in need of justification, and to claim that such beliefs are in fact justified because they are supported by my sensory evidence. For more on the Moorean approach, see G.E. Moore, "Proof of an External World," reprinted in his *Philosophical Papers* (London: Allen & Unwin, 1959), 126-148; Annalisa Coliva, "The Paradox of Moore's Proof of an External World," *The Philosophical Quarterly* 58 (2008): 234-243; Ernest Sosa, "Moore's Proof," in his *Reflective Knowledge* (Oxford: Oxford University Press, 2011), 3-23; Jessica Brown, "Doubt, Circularity and the Moorean Response to the Sceptic," *Philosophical Perspectives* 19 (2005): 1-14; and Susana Nuccetelli, "Sosa's Moore and the New Dogmatists," *Metaphilosophy* 40 (2009): 180-186. Nor is my response to the skeptic similar to the Russellian "inference to the best explanation" response (see Bertrand Russell, *The Problems of Philosophy* (Oxford: Oxford University Press, 1912/1997), Chapter II). For I am not saying that I can make an abductive inference to the conclusion that I am a brain in a body from some premises; rather, I am claiming that my belief that I am a brain in a body is justified in virtue of being supported by my evidence.

Reply: The skeptic is right that my evidence is consistent with BIVH. But the skeptic is wrong about whether my evidence supports the proposition that I am a brain in a vat just as much as it supports the proposition that I am a brain in a body. For my evidence supports the proposition that I am a brain in a body to a very high degree, and it does not support the proposition that I am a brain in a vat at all.

Why does my evidence support the proposition that I am a brain in a body? Because that is how my evidence makes it seem. My evidence makes it seem like I have hands, and arms, etc. (And also there is the Sally argument. Sally's evidence is misleading, which entails that it is leading somewhere. And my evidence is relevantly like Sally's.)

Why does my evidence not support the proposition that I am a brain in a vat? Well, here are some possible ways my evidence could support the proposition that I am a brain in a vat: (i) There could be glitches in the soundtrack of my life, as described in footnote 1 above. (ii) I could have glitchy visual sensations of hands, which would make it look like my hands were flickering in and out of existence. (iii) I could have odd, not-quite-right tactile sensations of having hands. Those are descriptions of possible evidence that would support the proposition that I am a brain in a vat. But none of those possibilities is remotely similar to what my evidence is actually like.

Here is a related point. The skeptic who says that my evidence supports N and BIVH equally well, because it is consistent with both propositions, is making a huge mistake about the relation of support between some evidence and a proposition. It is not (as the skeptic thinks) that some evidence equally supports any hypothesis that is consistent with that evidence. I.e., the following principle is false.

The Consistency Principle: For any propositions, p and q , and body of evidence, e , if p and q are both consistent with e , then e supports p and q equally well.

Here's an argument against The Consistency Principle. Suppose you're playing blackjack and have been counting the cards. Suppose you know that there are 10 cards left in the deck, and that 9 of them are face cards. This evidence supports the hypothesis that the next card will be a face card. The evidence is also consistent with the hypothesis that the next card will be a number card, but it does not support that hypothesis *to the same degree*. If you're a betting person, you should bet that the next card will be a face card.

Or suppose you meet some guy who performs an extraordinary act of kindness. Your evidence supports the proposition that this is a nice guy. Your evidence is also consistent with the proposition that this guy is literally the devil,

and has some nefarious reason for convincing you that he is a nice guy; but your evidence does not support this devil hypothesis *as strongly* as it supports the mundane hypothesis that the fellow is a nice guy.

Likewise, my sensory evidence is consistent with BIVH, but it does not support that hypothesis to the same degree that it supports the hypothesis that I have hands and a body. In fact, my evidence doesn't support BIVH at all; for my evidence doesn't make it *seem* like I am a handless, disembodied brain in a vat.

Objection 2: "The Consistency Principle is admittedly a bad principle. But there is something in the ballpark that is true, and that entails that your evidence supports BIVH as well as it supports N. For it's not just that your evidence is consistent with BIVH. *It's that this is exactly how things would seem to you if BIVH were true!*"¹⁵

Reply: The skeptic is right in claiming that this is how things would seem to me if I were a brain in a vat (provided the alien scientists were good at what they were doing, and chose to give me sensations like these). But the skeptic is wrong to claim that it follows that my evidence supports N and BIVH equally well. Why? Well, this is the very plausible principle that the skeptic is appealing to here.

The Counterfactual Principle: For any propositions, p and q , and body of evidence, e , if e is qualitatively identical to what the evidence would be like if p were true, and also qualitatively identical to what the evidence would be like if q were true, then e supports p and q equally well.

Although it is initially very plausible, The Counterfactual Principle is actually false. Here is a counterexample. Suppose I have seemings as of an empty car parked on the street. This is how things would seem to me if there were an empty car parked on the street. But it is also how things would seem to me if there were a car filled with a dozen cleverly hidden clowns. And yet my evidence does not support these two propositions equally well, since it supports the empty-car hypothesis to a much higher degree than it supports the clown-car hypothesis. After all, my evidence does not make it seem like there are a dozen cleverly hidden clowns in that car – instead, the evidence makes it seem like the car is empty.¹⁶

¹⁵ I am grateful to Katia Vavova for raising this objection.

¹⁶ If you are tempted to say that the evidence *does* support the clown-car hypothesis as much as it supports the empty-car hypothesis, consider that there are an infinite number of such crazy hypotheses (the monkey-car hypothesis, the unicorn-car hypothesis, etc.). The evidence cannot support each of them to the same non-zero degree.

The upshot is that since the objection is based on The Counterfactual Principle, and the principle is false, the objection fails.

Objection 3: “It is notable that we have these two hypotheses, BIVH and N, and each of them predicts that things will seem to you the way they do. In fact, the two hypotheses predict your evidence to exactly the same degree. So if the evidence supports one of these hypotheses, then it supports the other one to the same degree.”

Reply: This objection is based on the following principle.¹⁷

The Prediction Principle: For any hypotheses, $h1$ and $h2$, and observation, o , if $P(o/h1) = P(o/h2)$, then $S(h1/o) = S(h2/o)$ (i.e., if the probability of o given $h1$ equals the probability of o given $h2$, then o supports $h1$ and $h2$ equally well).

The Prediction Principle (like The Consistency Principle and The Counterfactual Principle) looks, at first glance, like a good principle, but it is not. Here are some counterexamples:

- The empty-car vs. clown-car hypotheses from above. Both predict that I will have experiences as of an empty car. But these empty car experiences I am having do not support the clown-car hypothesis as well as they support the empty-car hypothesis.
- The Obama-is-real vs. the Obama-is-an-elaborate-hoax hypotheses. Both hypotheses predict the experiences I have had as of there being a President Obama, but my experiences do not support the elaborate-hoax hypothesis to the same degree.
- The card counting case above. Let $H1$ = the hypothesis that 9 of the 10 remaining cards are face cards and the next card will be a face card. Let $H2$ = the hypothesis that 9 of the 10 remaining cards are face cards and the next card will be a number card. And let O = my current total evidence. Then $P(O/H1) = P(O/H2)$, but O does not support these two hypotheses equally well (since it supports $H1$ to degree .9 and $H2$ to degree .1).

These cases show that it is easy to come up with counterexamples to The Prediction Principle. Since the objection under consideration is based on that principle, and since the principle is false, the objection fails.

¹⁷ Why the shift from talk of propositions and evidence to talk of hypotheses and observations? Just because the objection sounds better when put in terms of hypotheses and observations. But I take it that hypotheses are a kind of proposition, and that observations are a form of evidence.

Objection 4: “The real problem is that your evidence doesn’t allow you to rule out BIVH. And since that hypothesis is a genuine skeptical alternative to N, your evidence doesn’t support N.”

Reply: The idea behind this objection is that you need to be able to rule out competing hypotheses before some hypothesis is supported by your evidence. So the objection is based on the following principle.

The Competing Hypotheses Principle: For any hypotheses, $h1$ and $h2$, and body of evidence, e , if $h1$ and $h2$ are competing explanations of e , and if e doesn’t allow the relevant agent to rule out $h1$, then e doesn’t support $h2$.

I admit that this is an initially very plausible principle. But it’s wrong. We have already seen several counterexamples: (i) the face card vs. number card hypotheses, (ii) the nice-guy vs. the devil hypotheses, (iii) the empty-car vs. clown-car hypotheses, (iv) the Obama-is-real vs. the Obama-is-an-elaborate-hoax hypotheses.

There is also a bigger problem with The Competing Hypotheses Principle. No matter what the evidence, and no matter what the hypothesis, there will always be some competing explanation (perhaps one that is undreamt of by the relevant agent) that is not ruled out by that evidence. So The Competing Hypotheses Principle will give the result that no evidence ever supports any hypothesis. And that can’t be right.

Objection 5: “It’s true that The Competing Hypotheses Principle is too strong, since it entails that no evidence ever supports any hypothesis. But there is something in the ballpark that is correct, and that doesn’t have the unwanted consequence that no evidence ever supports any hypothesis. If you are contemplating some hypothesis, and if there is some competing hypothesis that you cannot rule out, *and if that competing hypothesis is somehow made salient to you*, then your belief in the original hypothesis cannot be justified.”¹⁸

Reply: This objection is based on the following principle.¹⁹

¹⁸ I am grateful to David Pitt for pressing this objection.

¹⁹ I take it that the principle discussed in the text is equivalent to this one, which is more in keeping with the original Competing Hypotheses Principle: For any hypotheses, $h1$ and $h2$, body of evidence, e , and agent, S , if e is S ’s total evidence, if $h1$ and $h2$ are competing explanations of e , if e doesn’t allow the relevant agent to rule out $h1$, and if $h1$ has become salient for S as an explanation for e , then e doesn’t support $h2$.

The Alternative-Hypothesis-Made-Salient Principle: You cannot be justified in believing some hypothesis, *h1*, if there is a competing hypothesis, *h2*, such that *h2* is salient for you and you are not in a position to rule out *h2*.

Like the principles considered above, The Alternative-Hypothesis-Made-Salient Principle is initially plausible. But also like the above principles, it is false. Here are some counterexamples:

- You are wondering where your colleague Leila is. All of your evidence points to the hypothesis that she is in her office. Then for no special reason we put some random slips of paper into a hat and pull one out that says “MARS”. The degree of support for the hypothesis that Leila is in her office does not change, and you are still justified in believing that hypothesis.
- A variation on the parked car case from above. You are thinking the evidence supports the empty-car hypothesis. Then a fan of The Alternative-Hypothesis-Made-Salient Principle mischievously says, “Clown-car?” This does not change the degree of support for the empty-car hypothesis, or the fact that you are justified in believing it.
- A variation on the card counting case from above. You are about to bet that the next card will be a face card, when your friend says (for no special reason), “It could be that the next card will be a number card!” This does not change the degree of support for the face card hypothesis, which you are still justified in believing.
- We carefully inspect the room and come to believe that the only larger-than-a-bug, living organisms in the room are several humans. But then I say, just for fun, “What about the hypothesis that there are a bunch of mouse-sized, purple, Martian anthropologists who are super shy and who are sneaking around the room watching us while all the time avoiding detection?” You ask me whether I have any special reason to believe this hypothesis, and I truthfully answer, “No.” This conversation does not change the degree of support for your belief in the humans-only hypothesis, and your belief in that hypothesis continues to be justified.

It’s important to appreciate that in these examples, the only difference between before and after is that some competing hypothesis has become salient for the relevant believer. In particular, it is not the case that any new evidence has come in that undermines the evidence for the original hypothesis, or that functions as a defeater for whatever evidence supported that original hypothesis. And it seems clear that, in these examples, the mere fact that some alternative hypothesis becomes salient for the relevant agent does not change the degree of support by that agent’s evidence for the original hypothesis.

Objection 6: “You say that your evidence supports the belief that you are a brain in a body, because your sensations make it seem to you like you have hands. But your evidence equally supports the belief that you are a brain in a vat, since your sensations also make it seem to you like you are a brain in a vat being fed handish sensations. After all, we can imagine some weird person, Hilary, whose sensory evidence is relevantly like yours, but who interprets her evidence differently. Hilary notices the handish sensations she is having and concludes, ‘Hey, it looks like I’m a brain in a vat who is being fed handish sensations!’ Hilary’s belief that she is a brain in a vat being fed handish sensations is supported by her evidence. And since her evidence is relevantly like yours, your evidence also supports the proposition that you are a brain in a vat being fed handish sensations. So your evidence supports BIVH just as strongly as it supports N.”²⁰

Reply: The objection involves saying that Hilary’s belief that she is a brain in a vat being fed handish sensations is supported by her evidence. But I think it is not. I think that when Hilary interprets her evidence as supporting the belief that she is a brain in a vat, she is doing it wrong. She is simply mistaken about what her evidence makes it seem like to her. Her evidence does not make it seem to her like she is a brain in a vat who is being fed handish sensations; instead, her evidence makes it seem to her like she is a brain in a body, with hands.

It is important to distinguish between the two parts of Hilary’s conjunctive belief: (i) that she is a brain in a vat, and (ii) that she is being fed handish sensations. Her evidence does support the second part. But not the first part.

To see why, consider another person, Dilary, who is having sensations similar to Hilary’s (and mine). Dilary notices the handish sensations she is having and concludes, “Hey, it looks like I am a creature with tentacles who is being fed handish sensations!” Dilary is right that it looks to her like she is being fed handish sensations. But she is wrong in thinking *it looks to her* like she is a creature with tentacles. There is nothing tentacle-ish about Dilary’s sensations (recall that Dilary’s sensations are similar to mine and Hilary’s), just as there is nothing robot-arm-ish or flipper-ish about her sensations.

Or consider Stilary. He believes, for no good reason, that every human is constantly accompanied by an invisible, intangible daemon, who hovers above the human, sometimes protecting the human (on occasions when the human is not harmed) and sometimes negligently failing to protect the human (on occasions when the human is harmed). Stilary experiences sensations of you walking into a room and thinks, “Hey, it looks like I am being fed sensations of a human and a

²⁰ I am grateful to Katia Vavova and Maura Priest for raising this objection.

daemon entering the room!” When it comes to interpreting his evidence and determining which propositions are supported by that evidence, Stilary, like Dilary, is doing it wrong. For his sensations support the proposition that a human has entered the room, but they do not support the proposition that a daemon has entered the room.

Meanwhile, it’s the same with Hilary. Her sensations do not point toward her being a brain in a vat (even if they are consistent with her being a brain in a vat, and even if they are what her sensations would be like if she were a brain in a vat). Her sensations point toward her having hands. Her sensations point toward her having a body. (Recall the argument above about how Sally (who is relevantly like me, but is a brain in a vat) has evidence that is misleading: We can all agree that Sally’s evidence is misleading; and if it is misleading, then it is leading somewhere false, namely, to the proposition that she has hands and a body.)

One upshot of these considerations is that first person reports of how things appear to an agent are fallible. It is possible to be mistaken about what your evidence makes it seem like to you.²¹

Another upshot is that neither Hilary nor Dilary nor Stilary has sensations that support any skeptical hypothesis. Their sensations make it seem to them like they have hands and bodies. And for the same reason, my sensations do not support the skeptical hypothesis that I am a brain in a vat.

Objection 7: “But how do you know it’s *true* that you’re a normal brain in a body? You just assumed that.”

Reply: I plead guilty. I did just assume that I am a normal brain in a body. But consider this proposition:

T = the proposition that it’s *true* that I am a normal brain in a body.

There is no significant difference between T and N. They are equivalent. Thus, what I have said above ensures that the conditions for knowledge are satisfied with respect to T as well as N.

- ✓ I believe T.
- ✓ T is true.
- ✓ My belief that T is justified.

So that’s how I know it’s *true* that I’m a normal brain in a body.

²¹ There is of course another usage of ‘what x makes it seem like to S’ – the phrase can be used to mean *the belief that x causes S to have*. But this usage has nothing to do with justification.

Do You Know That You Are Not a Brain In a Vat?

Objection 8: “You say that you’re a normal brain in a body. But how do you *know* that? After all, for all you know, you might be a brain in a vat.”

Reply: Sometimes “How do you know that p?” is a question about your justification for believing p. But I have already answered that question. My justification is all of the evidence I have pointing toward my being a brain in a body. Other times “How do you know that p?” is a question about whether the philosophical conditions for knowledge have been satisfied in the case of your belief that p. But I have already answered that question, too.

Objection 9: “It feels like you’re cheating. You’re basically saying that IF you’re not a brain in a vat, THEN you can know lots of things about the external world. But the whole question is supposed to be whether you ARE a brain in a vat!”

Reply: No, that is not what the whole question is supposed to be. For I agree with the skeptic about the truth of the following two claims, which should both be completely uncontroversial.

Claim 1: Brains in vats don’t know much about the external world.

Claim 2: If I am a brain in a vat, then I don’t know much about the external world.

But I disagree with the skeptic about the truth of the following two claims.

Claim 3: Brains in bodies don’t know much about the external world.

Claim 4: If I am a brain in a body, then I don’t know much about the external world.

Since Claim 1 and Claim 2 are completely uncontroversial, the “whole question” is not about those claims. It is about Claim 3 and Claim 4. And my claim is that, since brains in bodies are justified in believing, for example, that they have hands, they know lots of things about the external world. So I’m not cheating.

Here is an analogy. Suppose I am dealing with a different kind of skeptic, a more mundane one, who says that I do not know where my car is parked (since it could have been stolen or towed since I last saw it). Suppose that we discuss this matter as we walk toward the spot where I think my car is parked. When we arrive, and see that my car is indeed there in the spot where I thought it was, my skeptical friend should not give up and say, “Oh, I guess you did know where your car was parked all along.” Instead, if she is a competent skeptic and a good philosopher, she should say, “Yay! I’m glad your car is where you thought it was,

Ned Markosian

so that we can now drive to the restaurant. But you still didn't know it was here a few blocks ago, when we started this conversation." This shows that it wasn't the *truth* of the relevant belief that was in question all along; it was whether the belief was *justified*.

It's the same with the controversy between the external-world skeptic and me. It is not the truth of the proposition that I am a brain in a body that is in dispute. It is whether that belief is justified. This is why it's not cheating to assume, as I have, that I am a brain in a body.

Objection 10: It might appear that I am committed to an inconsistent triad. First, I have said that the *degree of support* that my evidence gives to the proposition that I am a normal brain in a body is very high (greater than .9, say, on a scale from 0 to 1). In short (using some natural notation),

$$(a) \quad S(N/e) > .9,$$

where 'N' is the proposition that I am a normal brain in a body, and 'e' is my total evidence.

Second, it's very natural to say that the *probability* that I am a brain in a body is either unknown or unknowable or inscrutable or undefined. (After all, who knows what the percentage of brains in vats is, throughout the universe, among brains having experiences like mine?) In short (using some familiar notation),

$$(b) \quad P(N/e) \text{ is unknown or unknowable or inscrutable or undefined.}$$

And third, it also looks like I will be committed to what we might call The Equivalence Thesis: The degree of support for some proposition by someone's evidence is equal to the conditional probability of that proposition given that evidence. In short,

$$(ET) \quad S(p/e) = P(p/e).$$

But (a), (b), and (ET) form an inconsistent triad, because .9 is not unknown or unknowable.

Reply: I think there might be independent reasons to question (ET), but it is certainly pretty appealing. In any case, I am not going to deny it here. Instead, I think that the best way to deal with this problem is to deny (b). My evidence (namely, my current sensations) not only makes it *seem* like I am a normal brain in a body, but also makes that proposition *very probable*.

Here is an argument for the claim that my current total evidence makes it very probable that I am a brain in a body. Suppose I get some very credible new evidence suggesting that the probability that I am a brain in a vat is .5. For example, suppose an oracle appears before me, convinces me that she is omniscient and reliable, and informs me that half of the brains like mine in the universe are brains in vats. This would be a life-changing experience that would seriously undermine my previous sensory experience for N. After the experience with the oracle, the probability that I am a brain in a body, given my total evidence, would be .5. But this would be because the new evidence from the oracle had *dramatically lowered* the probability of N. Which suggests that the probability of N given my evidence before the oracle showed up was *extremely high*. In fact, it looks like the skeptic pushing this objection has to maintain that the new evidence from the oracle *doesn't change the probability of N at all*, and that seems crazy.

Now, you might think the skeptic can get around this argument by saying that the new evidence from the oracle changes the probability of N, given my total evidence, from unknown to .5. But I don't think that response on the part of the skeptic can work, because in other, less dramatic cases – like a brief glitch in the sensory track, or an encounter with a less convincing oracle – it is clear that the new evidence is lowering the probability conferred by my total evidence *incrementally*, rather than flipping it in a binary way from unknown to .5 or whatever.

Objection 11: “Given your evidence, the probability that you are just a brain in a vat is equal to the probability that you are a brain in a body. It's 50-50. And if two propositions are equally probable, given some evidence, then that evidence supports them equally well. So your evidence supports BIVH just as much as it supports N.”

Reply: This objection involves an appeal to the following principle.

The Probability Principle: For any propositions, p and q , and body of evidence, e , if $P(p/e) = P(q/e)$, then $S(p/e) = S(q/e)$.

The Probability Principle, unlike some of the other principles we have seen, is a good principle. Or at any rate, it follows from (ET), so it is as good as (ET), which I am accepting for the purposes of this paper. The problem with the objection is the claim that $P(\text{BIVH}/e) = P(N/e)$ (where e is my total evidence). For as I have just argued (using the oracle example), $P(N/e)$ is very high, while $P(\text{BIVH}/e)$ is super low.

Here is a bonus argument for this latter claim. It seems clear that $S(N/e)$ is very high, and also that $S(BIVH/e)$ is very low. (I have given two arguments for this conclusion. The first was an argument from the claim that my sensations make it seem like I am a brain in a body, not a brain in a vat. And the second argument for this conclusion was from the claim that Sally (who is an unfortunate brain in a vat) has evidence that is misleading.) So by (ET) we immediately get the result that $P(N/e)$ is high, while $P(BIVH/e)$ is low.

Notice that if the skeptic tries to resist this argument by rejecting (ET), then the objection currently under consideration cannot get off the ground, since it needs (ET) to justify The Probability Principle.

5. A Final Comment on the Problem of External-World Skepticism

I want to make one final comment on the problem of external-world skepticism. You might have a lingering feeling that I have not really solved the problem. If so, I think this is likely due to a confusion of two different problems.

First, there is an *epistemological problem*: How can I respond to the skeptic who argues that I don't know anything about the external world? To solve this problem I would need a reason to reject one of the premises of The Brain In a Vat Argument. And as we have seen, if I am to reject premise (1) of that argument, this will come down to showing that most of my ordinary beliefs about the external world, including my belief that BIVH is false, are justified.

Second, there is a *practical problem*: Is there some test or procedure that I can perform to determine which situation I am in – a brain in a vat situation or a brain in a body situation? To solve this problem I would need to come up with something analogous to the pinch test for dreaming. That is, I would need to come up with a “pinch test” for the brain in a vat scenario, such that if I perform this test, then it will all of a sudden be revealed to me whether I am a brain in a vat.

I claim to have solved the first problem, the epistemological one. I have argued that ordinary beliefs about the external world, including my belief that BIVH is false, are justified by our sensory evidence. Which means that ordinary knowledge claims about the external world are true – including my belief that I am not a brain in a vat. But I do not claim to have solved the second problem. In fact, I think there is no solution to this practical problem. The answer to the question is *No, there is no “pinch test” for the brain in a vat scenario*. But luckily, this does not change the fact that I do know lots of things about the external

Do You Know That You Are Not a Brain In a Vat?

world, contrary to what the skeptic and The Brain In a Vat Argument have claimed.²²

²² Earlier versions of this paper were presented at Western Washington University, The University of Wisconsin – Milwaukee, Cal State – LA, UC – Irvine, Boğaziçi University, Yerevan State University, The University of Illinois, and The University of Alabama. I am grateful to members of all eight audiences for helpful comments. I'm also grateful to Katia Vavova and Dan Korman for helpful comments on an earlier version of this paper, and to Sharon Ryan for helpful discussions of this topic.