

How Can We Know that We're Not Brains in Vats?

Keith DeRose
Yale University

1. Do We Know that We're Not BIVs?: The Threat of Skepticism and "Moorean" Responses to It

Consider the hypothesis that you are a bodiless brain-in-a-vat who has been electrochemically stimulated to have precisely those sensory experiences you have had, perhaps because you are appropriately hooked up to an immensely powerful computer, which, taking into account the "output" of the brain that is you, has seen to it that you receive appropriate sensory "input" (henceforth, a "BIV"). Do you know that you're not a BIV?

Often in epistemology, we ask *how* we know various things which are such that it's pretty clear to almost everybody that we *do* indeed know them—though perhaps some crazed skeptics may deny that we have the knowledge in question. That's not the case here. Intuitions vary widely. Indeed, some epistemologists seem to think it's pretty clear that we *don't* know that we're not BIVs, and not because they have arguments for this position, but because it seems to them intuitively clear enough to function as a starting point for inquiry. I've spoken with many philosophers, however, who are puzzled by this stance, either because the issue seems far from clear to them, or because they're actually inclined to give the opposite answer—that we *do* know we're not BIVs. If you ask nonphilosophers

Keith DeRose is Associate Professor of Philosophy at Yale University. His primary areas of research are epistemology, philosophy of language, and history of modern philosophy. Recent articles include "Contextualism: An Explanation and Defense," in The Blackwell Guide to Epistemology, edited by J. Greco and E. Sosa; "Conditional Assertions and 'Bisquit' Conditionals," with Richard Grandy, Nous, 1999; and "Can It Be That It Would Have Been Even Though It Might Not Have Been?" Philosophical Perspectives, 1999. Along with Ted Warfield, he edited an anthology of recent work on Skepticism for Oxford University Press (1999).

(after explaining the BIV hypothesis) whether they know they're not BIVs, I think you'll find both positive and negative answers well represented.¹

The pressure for saying that we know we're not BIVs, then, doesn't come so much from it's being so intuitively clear in itself that we know this, but from the threat of skepticism. Though it's not so clear in itself that we know we're not BIVs, it does seem quite clear—to most of us, anyway—that we know such things as that we have hands, and the admission that one doesn't know that one is not a BIV is one of the key premises in one of the most powerful skeptical arguments to the conclusion that one doesn't even know that one has hands. This argument in its most basic form can be rendered as follows, where O is some proposition one would ordinarily think one knows (e.g., I have hands) and H is a suitably chosen skeptical hypothesis (e.g., I am a BIV):

The Argument by Skeptical Hypothesis:

1. I don't know that not-H.
 2. If I don't know that not-H, then I don't know that O.
- So,
C. I don't know that O.

To admit that one doesn't know that one isn't a BIV, then, is to accept the first premise of this skeptical argument, and to be left with two options: either skepticism—accepting the conclusion that, for example, one doesn't know that one has hands—or denying the second premise of the skeptical argument. In the present paper, I will just note the existence of these two options, each of which has been embraced by prominent epistemologists, and will not argue against them. I find them both quite unattractive and unpromising, even after studying the efforts that have been expended on their behalf. My topic here is what we can call "Moorean" responses to the skeptical argument: approaches according to which we do know that we're not BIVs.² But it will be important in discussing such approaches to keep in mind that the motivation—or at least a main motivation—for thinking that we know that we're not BIVs is to avoid skepticism—or, more exactly, to avoid the choice between skepticism and another option that many find quite unpalatable (denying the second premise of the skeptical argument).

Judging by the amount of literature it has generated in the last 18 years in philosophical journals, a very prominent answer to my title question ("How can we know that we're not BIVs?") is: Via a Putnam-style argument from Semantic Externalism that is, in an appropriate sense, a priori. In calling this answer "prominent," I don't mean to suggest it is very widely accepted. In fact, a good portion—probably most—

of the literature on this approach has been critical. Here, I will continue the attack against this answer to our question, but I will also point out a limited way in which this approach may have some merit.

In recent years, a different type of Moorean approach has emerged—contextualist Mooreanism. This is the approach I advocate. I will not here attempt a full-scale defense of this alternative approach. However, after I argue that the Putnam-style answer to our question suffers from various defects, I will explain how the alternative approach I favor holds promise for giving us an answer that avoids those defects.

One might be interested in the questions of whether, and, if so, how we know that we're not BIVs either to investigate whether BIV-inspired skepticism can be cut off via a denial of the first premise of the skeptical argument, or one might be interested in such questions in their own right, regardless of whether such knowledge can be used to combat skeptical arguments. In sections 2–7, we'll consider accounts of how we might know we're not BIVs as responses to skepticism. I will argue that Putnamian responses from Semantic Externalism to BIV-inspired skepticism are highly unsuccessful. Indeed, as a response to skepticism, the "Putnamian" response seems hopeless—as I suspect Putnam himself realized. But we'll next turn, in sections 8–9, to the question of whether Putnam-style inferences might provide a priori knowledge of the fact that we're not BIVs, even if they can't thereby provide relief from skepticism. It's here that the approach seems to me to hold some promise.

2. The Arguments from Semantic Externalism

Mooreans often face our title question as a very pointed question. In fact, after claiming that we do know that we're not BIVs, the opening line of attack one is most likely to face simply consists of being asked: *How could we possibly know such a thing?*

Well, what if we could *prove* that we're not BIVs? That, it at least seems, would provide the boldest, strongest, and most definitive possible response to the pointed question. Famously, Hilary Putnam has claimed to have "an argument we can give that shows we are not brains in a vat."³ Now, Putnam himself seems largely uninterested in using this argument to combat skepticism,⁴ but, understandably enough, others have seized on Putnam's argument, or closely related arguments to the conclusion that we're not BIVs, as antiskeptical weapons.

Such "Putnamian" arguments are grounded in Semantic Externalism, according to which the contents of at least some of

one's thoughts are not completely determined by "internal" facts about what is going on inside one's head, but are at least partially determined by such "external" facts as the nature of the items one has been in contact with. In particular, according to Putnam, you cannot think about, say, trees, if you haven't been causally connected with trees in the proper way. Thus, a BIV, since it hasn't been in contact in the proper way with real trees, cannot refer to or think about trees. When such a BIV thinks such thoughts as those it expresses via the sentences "There's a tree," or "Here's a hand," or "I'm not a BIV," then, it is not thinking the same thing that those words would express in our mouths/minds (given that we're not BIVs). What does "tree," as used by a BIV (in "vat-English"), refer to? Putnam lists several different suggestions: "It might refer to trees in the image, or to the electronic impulses that cause tree experiences, or to the features of the program that are responsible for those electronic impulses." All of these suggestions are in the spirit of Semantic Externalism, because, as Putnam writes, "there is a close causal connection between the use of the word 'tree' in vat-English" and each of these suggested referents.⁵ Importantly, on any of these suggestions, the BIV ends up thinking something true when it thinks "There's a tree," or "Here's a hand," or even "I'm not a BIV," for, to take the "in the image" reading, the BIV is indeed indicating a tree-in-the-image and a hand-in-the-image, and it indeed is not a BIV-in-the-image (it's just a BIV).

But how might such an argument proceed from Semantic Externalism to the conclusion that one is not a BIV? Two quite different ways have been proposed, both as promising anti-skeptical strategies in their own right and as interpretations of Putnam.⁶ Since I will pursue criticisms of the arguments that don't depend on the fine points of the arguments, we needn't go into great detail about them. Still, we need a brief, sketchy statement of them to begin.

The main idea of the first type of argument—which we'll call the "Dilemma argument"—is this. If I am a BIV, then by "I am not a BIV" I mean that I am not a BIV-in-the-image (or some closely related true thing), which is in that case true. On the other hand, if I am not a BIV, then by "I am not a BIV," I mean that I am not a BIV, which is in that case true. Thus, whether I am a BIV or whether I am not, my use of "I am not a BIV" is true. Either way, it's true; so, it's true: I'm not a BIV.⁷

We'll call the second type of argument the "Compatibilist argument."⁸ It combines a negative externalist claim about what a BIV does *not* (or cannot) mean or think—that by "tree," "hand," "vat," etc., the BIV does not refer to trees, hands, vats, etc.—with a positive claim to the effect that we do have the thoughts in question—the thoughts that the BIVs cannot have. These together imply that we are not BIVs.

3. Old Objections: Varieties of Semantic Externalism and Varieties of Skeptical Hypotheses

To set up my main criticism of these arguments, it will be helpful to first review and then assess a couple of the objections that have already been leveled against them in the literature, and which are by now well-known—oldies, but goodies.

The first of these objections is particularly damaging against the first, "Dilemma" form of the argument, though, as we'll see, it can combine with the second objection we'll look at to also target the Compatibilist form of the argument. The objection is that it's far from clear that a form of Semantic Externalism that's strong enough to fuel the antiskeptical argument is correct.

There *is* strong support for the claim that the contents of at least some of one's thoughts are not completely determined by "internal" facts about what is going on inside one's head, but are at least partially determined by such "external" facts as the nature of the items one has been in causal contact with. This support comes mainly from thought experiments—most famously, those put forward by Putnam and Tyler Burge⁹—in which it's clear, at least to most who consider them, that thought contents can differ due to "external" factors. But this "low-grade" externalism (by which I mean the bare claim that thought contents can differ due to differences in "external" factors), by itself, cannot fuel the antiskeptical arguments we're considering. What does this claim tell us about the content of the thoughts of BIVs? By itself, not much. It alerts us to the *possibility* that the contents of the thoughts of BIVs are different from the contents of our own thoughts (given that we're not BIVs). Since the nature of BIVs' contact with reality is so radically different from our own, we may well suspect it likely that their relevant thoughts differ in content from the analogous thoughts we have, but without further information about just when and how external differences produce differences in content, "low-grade" Semantic Externalism doesn't tell us what those BIVs are thinking. The thought experiments used to support Semantic Externalism, of course, may do more than simply establish the "low-grade" claim; they may give us some idea of when and how the contents of thoughts are sensitive to "external" factors. But do they establish a strong enough brand of externalism to fuel a Putnam-style argument to the conclusion that one is not a BIV?

The Dilemma form of the antiskeptical argument requires the claim that when a BIV thinks the words "I am not a BIV," it thinks something true. For our purposes, let's call forms of

Semantic Externalism that imply this "high-grade" externalism. While the thought-experiments used to support Semantic Externalism often elicit strong intuitions, it is *far* from intuitively clear that BIVs think something true when they think "I am not a BIV," and this case is different enough from the standard thought-experiments, like the Twin Earth cases, for there to be any clear path from the relatively clear intuitive verdicts about the standard cases to the intuitively cloudy claim that BIVs think something true. It's worth noting here that Putnam himself seems quite unsure that the BIVs succeed in meaning *anything at all*, he writes that if we are BIVs, then "what we now mean by 'we are not brains in a vat' is that we are not brains in a vat in the image or something of the kind (*if we mean anything at all*)," and, along the same lines, "So, if we are brains in a vat, then the sentence 'We are brains in a vat' says something false (*if it says anything*)."¹⁰

Indeed, one of the main apparent advantages of the Compatibilist form of the argument over the Dilemma form is that the Compatibilist argument, initially, at least, does not appear to require such high-grade externalism. Rather than needing any claim to the effect that the content of BIVs' thoughts is such that they express truths by various sentences, it requires only negative claims to the effect that BIVs can't have various thoughts that we have. So, for instance, the argument might go like this:

Compatibilist Argument

1. I think that water is a liquid.
 2. No BIV can think that water is a liquid.
- So,
C. I am not a BIV.

Premise 2, above, is in several respects a fairly weak externalist claim. If *any* of Putnam's positive proposals as to what the contents of BIVs' thought might be—whether by "water" BIVs refer to water-in-the-image, or to the electronic impulses that cause water experiences, or to the features of the program that are responsible for those electronic impulses—2 is on solid ground, for on none of these proposals do BIVs refer to water by "water." What's more, if Putnam's worry that BIVs may not succeed in meaning anything at all should prove to be well-founded, then 2 is still right. Indeed, 2 is true on any account of what the BIVs mean by "water" other than that they mean water. Though the "mid-grade" externalist claim that they don't mean water does go beyond the bare "low-grade" claim that external differences can result in differences in thought content, it seems not to go beyond it in any way that should produce much doubt in anyone who goes along with the spirit of Semantic Externalism.

At least this seems so *if* we understand the BIV hypothesis in certain ways. And here we come to the second, often noted,¹¹ problem for the antiskeptical strategy: that it works only against some, but not all, of the ways the BIV hypothesis can be construed. To solidify his externalist claim that the BIVs he was imagining were not capable of thinking about trees, hands, vats, etc., Putnam imagined a very special scenario in which the BIVs have always been BIVs. In fact, he went further and supposed that all sentient beings had always been BIVs, the universe, by accident, just happening "to consist of automatic machinery tending a vat full of brains."¹² But what of other scenarios? Putnam's version of the BIV hypothesis seems to be, in David Christensen's apt words, "cooked up to be vulnerable to the semantical reply."¹³ If instead we imagine that the computer tending the brains were programmed by fully-bodied humans who have experienced real trees, water, etc., then even some externalists might start to wonder whether the BIVs might not succeed after all in referring to water by their use of "water," through their indirect contact with water that goes through the programmers. If we go further, as any smart skeptic should, and try to "cook" the scenario in order to make it less vulnerable to the antiskeptical reply, we can, as has been suggested in the literature (see note 11), construe the hypothesis so that I am a BIV who has only recently been envatted, after many years of normal embodiment and causal contact with real trees, hands, vats, etc. Though it's quite unnecessary, we can go even further, and imagine a scenario of *very* recent and *very* temporary envatment: Not only was I normally embodied until just recently, but soon I will be returned to a normally embodied state, remaining all the while oblivious to these drastic changes in my situation. Then, even to most externalists, it seems that, even now, while envatted, I do mean tree by "tree," water by "water," vat by "vat," etc., and so I can think that water is a liquid and would be thinking something false when I think "I'm not a BIV."

This is a nasty problem for Putnam-style antiskeptical arguments of both varieties. Ted Warfield, in his presentation of his version of the compatibilist argument, seeks, if somewhat halfheartedly, to defend the argument from this objection as follows:

While I admit that there is a certain intuitive force to this objection to the scope of semantic antiskeptical arguments, I would like to suggest in closing that we should not be too quick to conclude that the range of applicability of semantic antiskeptical arguments is extremely narrow. After all, just which skeptical hypotheses are and are not vulnerable to such arguments is a function of the *details* of the externalist necessary conditions on thought and reference. We may have

certain suspicions about what these details will look like...., but I do not think that we know enough about the semantics of thought and reference to be very confident in the truth of such suspicions.

The range of applicability of semantic antiskeptical arguments depends on how these sorts of issues are resolved and I see no way of resolving them without serious and detailed work on psychosemantic theories. The common deflationary view of Putnamian antiskepticism is therefore at least a bit premature.¹⁴

Warfield is right about several things above. In particular, he's right about how the details of the externalist account of content will interact with the exact nature of the hypothesis in question to determine whether, as far as these objections go, the hypothesis is susceptible to the antiskeptical argument from externalism. And he's right that it's far from clear just what range of hypotheses will fall prey to the strategy.

But, for the purposes of providing relief from *BIV-inspired skepticism*, so what? Perhaps, as Warfield suggests, the range of BIV hypotheses susceptible to the strategy may not be "extremely narrow." Perhaps, at the end of the day (when we get the details of externalism figured out), we'll all be surprised at how many versions of the BIV hypothesis can be shot down Putnam-style. Let's be very generous to the antiskeptical strategy and actually suppose that, at the end of the day, the range of skeptical hypotheses that *escape* being shot down is extremely narrow. *Still the skeptic wins*: She just needs one version of the hypothesis to work to establish her skeptical conclusion that we don't know such things as that we have hands.¹⁵

Putnam seems to have been wise not to try to promote his argument as an antiskeptical weapon!

4. Objections: The Disadvantages of Heroism

Things get worse. Of course, as those who have followed the literature know, there are other potential problems for the antiskeptical arguments we're investigating. Some of these appear to me quite threatening. But rather than rehearsing them, I want instead to explore another line of objection that is more fundamental in that, in addition to amplifying the effect of the other problems to the extent that they are or *even appear to us to be* real problems, also calls into question the antiskeptical value the Putnam-style arguments would have *even if they did work*.¹⁶

The Putnam-style response to skepticism attempts—leaving open the question of whether it succeeds—to provide an

argument that one is not a BIV. In this way, it seems a very aggressive antiskeptical strategy. When a skeptic challenges your presumed knowledge that, say, you have hands, by suggesting that you just might be a BIV, what bolder answer could there be than to prove to her that you're no such thing? But in another respect, this strategy concedes much to the skeptic, and therefore provides a very weak response. Too weak for me. By proving that one is not a BIV, one seeks, it seems, in following this strategy, to *make* it the case that the first premise of the Argument by Skeptical Hypothesis is false, as applied to oneself. In this respect, this strategy is "heroic" in the way Descartes's response to the evil genius argument is: The Putnam-style arguer, like Descartes, seeks by constructing a proof against the obtaining of the relevant skeptical hypothesis, to *gain* knowledge, for himself and all that would follow him, that the hypothesis is false. Externalist semantics has replaced Descartes's God as the slayer of skeptical hypotheses. Presumably, though, in either case, the proof only helps those who follow the hero—who know and understand the argument. But what of people who have never encountered this sophisticated argument that one is not a BIV? For instance: What about my Mom?! Does *she* fail to know that she has hands?¹⁷ Nonheroic strategies, which attempt not to show how to gain knowledge in the face of the skeptical argument, but rather to show how the skeptical argument never worked in the first place, by protecting the knowledge of the unphilosophical, seem in that important way to be more aggressively anti-skeptical.

Of course, one can give an argument for a conclusion one already knows to be the case. The Putnamian arguer can, therefore, object to my describing him as seeking to make it the case that he knows that he's not a BIV by means of his argument. The fact that he's providing an argument for the conclusion that he's not a BIV doesn't commit him to the position that he didn't already know this and that others, who have no proof against the hypothesis, don't know this.

Fair enough. But we want to know: How then *does* this argument that I am not a BIV help me with respect to the skeptical argument? The skeptic, recall, argues as follows:

The Argument by Skeptical Hypothesis:

1. I don't know that not-H.
2. If I don't know that not-H, then I don't know that O.

So,

- C. I don't know that O.

The argument is clearly valid, and each of its premises is at least fairly plausible, so it at least threatens to show that I don't know that O. How does a Putnamian argument to the

conclusion that I'm not a BIV help in this predicament? That there is such an "antiskeptical" argument, even should it prove sound, doesn't seem to have any tendency to show that the skeptical argument is invalid, or that its second premise is false, or that its first premise wasn't true of me before I knew of the antiskeptical argument. The only *readily apparent* way the antiskeptical argument could help is by *giving* me knowledge, should I not already have it, that not-H, and thereby making the first premise false, as applied to me, if it isn't already false. Those who discuss the semantic antiskeptical strategy—both those who defend it and those who attack it—tend not to explain how the presence of an argument to the conclusion that one is not a BIV might help us out of our skeptical predicament in any other, not-so-readily apparent way.¹⁸

I'll leave it to the promoters of the semantic antiskeptical strategy to tell us how it's supposed to block the skeptical argument if not by heroically giving us knowledge that we're not BIVs by means of our grasp of this argument and thereby making the first premise of the skeptical argument false (if it isn't already false).¹⁹ In the meantime, I'll assume that that's how the semantic argument is supposed to help. If that *is* how the antiskeptical argument helps, then, of course, as I've charged, it doesn't help at all in protecting the knowledge of the nonphilosophical who have never grasped the argument. The problem doesn't stop there. What's more, even among the philosophical, of course, many who study the Putnam-style arguments, for various reasons (a couple of which we looked at), don't even believe (much less know) that they work, and the knowledge of these philosophers is not protected by the strategy. Even among those who promote this antiskeptical strategy, I haven't met a single one who actually believes the strategy works against *all* the versions of the BIV hypothesis it would need to work against in order to protect one's knowledge that one has hands.

Here, the problems inherent in the "heroic" nature of this response start to interact in very destructive ways with the other problems that we looked at earlier. Note, for instance, how some of the considerations Warfield puts forward in *defense* of the Putnam-style strategy—that we don't yet know enough about the details of thought and reference to be confident about what range of hypotheses it will work against—start to look like *attacks* on the strategy in light of the realization that the strategy only protects our knowledge insofar as we do know that the skeptical hypotheses are false by means of the Putnam-style arguments.

If, as we saw earlier, the externalist strategy must work against even recent envatment scenarios in order to provide the needed strategy that will only work for those who grasp the externalist argument, there is a very substantial chance,

in my estimation, that *even if the antiskeptical argument against recent envatment is sound* (which it surely isn't), there is not a single human on the face of the Earth whose knowledge is protected by it from the skeptical argument. Has anybody actually come to *know* that they're not a recently envatted BIV by means of a Putnam-style argument? Even if such an argument for the falsity of recent-BIV hypotheses were sound, could someone gain knowledge of that falsity by means of an argument whose soundness is doubted, if not outright denied, by even those externalists who have studied the argument most carefully? Perhaps. But I think it's fair to say that the range of people who have come to know that they're not recent BIVs by means of such an argument is, at most, exceedingly narrow. But then, since, as we're now assuming, the argument protects one's knowledge from the skeptical argument only by giving one knowledge that the skeptical hypothesis is false, and thereby making the first premise of the skeptical argument false, as applied to one, then the range of people whose knowledge is protected by means of this antiskeptical argument is extremely narrow.

Of course, for all I've argued, if the Putnam-style arguments are in fact sound, and sound as applied to all the hypotheses they'd have to work against, then there is some (but not much) hope that some day in the distant future, as their soundness comes to be known and then widely recognized and even widely known, they may come to give many knowledge that they're not recent BIVs, and thereby come to block the skeptical argument, as it's applied to many. But if you're worried that the skeptical argument shows that right now you don't know that you have hands because you don't know that you're not a BIV, then there doesn't seem to be any relief to this worry provided by the Putnam-style response.

One final worry concerning the heroic nature of this response: I, for one, take my knowledge of the fact that I have hands to be a very solid piece of knowledge indeed—warranted far above and beyond the call of knowledge.²⁰ But if, as those who accept premise 2 of the skeptic's argument would seem to agree, my knowledge that I have hands depends on my knowing that I'm not a BIV, then there's room to worry that even if I do know that I'm not a BIV, but just barely know it (if my belief that I'm not a BIV is warranted to a degree sufficient for knowledge, but not to any much higher degree), then, though my knowledge that I have hands is preserved, it is not preserved as the piece of highly warranted knowledge that I take it to be. Though I find the externalist thought-experiments to be pretty convincing, *so far as philosophy goes*, I have my doubts that, based on them, I have enough warrant to *know* at all (much less to know in any especially solid manner) that even low-grade externalism (much less the stronger forms of

externalism that are needed to fuel the semantic arguments) is true. Now, I don't think there's any clear and obvious path from the premise that one doesn't know that O if one doesn't know that not-H to the conclusion that one's knowledge that O cannot exceed one's knowledge that not-H in the degree of its warrant. Still, there is room to worry that whatever knowledge that not-H can be gained for us via the semantic argument will be too shaky to support the kind of highly warranted knowledge we want of such facts as that we have hands. And, at any rate, we should recognize that, in order for the semantic reply to skepticism to work against a skeptical hypothesis, the semantic argument involved has to be strong enough and clear enough that we can, on its basis, come to *know* that the hypothesis is false.

5. What Would a Nonheroic Alternative Look Like?

We need a nonheroic response to the skeptical argument: one that seeks not to gain or regain knowledge of such things as the existence of our hands in the face of the skeptical argument, but one according to which the skeptical argument never worked in the first place—one according to which it was never the case that both of the skeptical argument's premises were true of us. Recall, however, that here we have set aside responses that deny the second premise of the skeptical argument.²¹

Any response to skepticism that accepts the skeptic's second premise faces a problem. Of course, there's the obvious problem that the skeptic's first premise enjoys intuitive support as well. But even for those who find the skeptic's first premise intuitively very questionable, and *even* for those whose dominant intuition is to reject that first premise, there is this problem: Since there is, for almost everyone, at least *some* substantial intuitive support for the skeptic's claim that we don't know we're not BIVs, even those who find this intuitive support outweighed by an opposing intuitive push toward saying we do know that we're not BIVs are apt to feel that this was a close call and that whatever knowledge we might have that we're not BIVs just barely squeaks by as a case of knowledge. Our epistemic position with respect to our not being BIVs can seem quite shaky: Even if it is good enough to make us knowers, it makes us shaky, just-barely knowers. This is troublesome if one has admitted that one's knowledge that one has hands—which we all take to be a pretty solid piece of knowledge—depends on one's knowing that one is not a BIV. Here the temptation to become heroic is great: Since our initial, prephilosophical position with respect to our not being BIVs seems shaky, it's natural to try to improve our position by

means of a philosophical argument. But we've seen the problems of heroism.

Nonheros who accept the skeptic's second premise seem to have three options: the good, the bad, and the ugly. We'll take them in reverse order. First, they can accept that our knowledge of such propositions as that we have hands is shaky, just-barely knowledge. This seems unacceptable. Second, they can seek to show that even though one cannot know that one has hands if one does not know that one is not a BIV, still, one's knowledge that one has hands can be about as firm and solid as we think it is even though one's knowledge that one is not a BIV is quite shaky. This strikes me as quite unpromising. As I wrote above, the admission that one doesn't know that O if one doesn't know that not-H doesn't clearly imply that one's knowledge that O cannot exceed one's knowledge that not-H in the degree of its warrant. (I don't think the implication holds at all—whether clearly or unclearly.) Still, in this case, I don't think that one's knowledge that one has hands can exceed one's knowledge that one is not a BIV in degree of warrant *by enough* to secure for the former the solidity it seems to have while the latter is as shaky as it seems to be, if the former piece of knowledge depends on the latter.

This leaves us with the third (or first, since we're taking them in reverse order) option: To hold that one's knowledge that one is not a BIV is not just a piece of knowledge, but a quite solid, highly warranted piece of knowledge. Recall, this is to be a nonheroic response, so we do not seek to shore up our knowledge that we're not BIVs, or to *make* it the case that we know this with a high degree of warrant, but we seek to defend that our prephilosophical, initial state is that of knowing quite solidly that we're not BIVs.

But how could this be? Why would our alleged "knowledge" that we're not BIVs strike us as so shaky—to the extent that, for many, it seems not to be knowledge at all—if in fact it is such a solid, highly warranted piece of knowledge? (How does this get to be the "good" option?) The appearance of shakiness, it seems, must be explained away as an illusion.

6. A Nonheroic Alternative: A Moorean Contextualist Account

I endorse (see my 1995) what can be profitably classified as a *contextualist Moorean* response to the argument from skeptical hypothesis, and in doing so, I take myself to be in at least rough alliance with Gail Stine, Stewart Cohen, and David Lewis.²²

We follow the basic contextualist strategy²³. We accept a contextualist theory of knowledge attributions;²⁴ we accept that

at some very unusually high standards for knowledge (which we'll here call the "absolute" standards), we don't count as knowing that we have hands; we claim that we do know that we have hands according to the much lower standards for knowledge that typically govern most of our ordinary conversations; and we seek to explain the persuasiveness of the skeptic's argument, at least in part, by claiming that the presentation of the skeptic's argument has at least some tendency to put into play the very "absolute" standards at which we don't count as knowing that we have hands.

In light of our concession to the skeptic that we don't know that we have hands according to the absolute standards for knowledge, ours is not a thoroughly antiskeptical response. The sense in which ours is an antiskeptical response is that we do seek to protect the claim that we do know such things according to many nonabsolute standards, and thus seek to protect the truth values of our ordinary claims to know.²⁵ In the way that ours is an antiskeptical response, ours is a Moorean response, because, according to the four of us, at the nonabsolute standards at which we know we have hands, we also know that we are not BIVs. More specifically, ours is a response according to which:

- (a) Premise 2 of the skeptical argument is true at whatever epistemic standard it is evaluated at.
- (b) Premise 1 and the conclusion of the skeptical argument are true when evaluated according to the unusually high, "absolute" standards for knowledge that the presentation of the skeptical argument has at least some tendency to put into play.
- (c) Premise 1 and the conclusion of the skeptical argument are false when evaluated at the standards for knowledge that are set by most ordinary contexts.

(By replacing all instances of "Premise 1" with "Premise 2," and vice versa, one gets a description of what might be thought of as a "contextualist Nozickian" response to the skeptical argument: a response that follows the basic contextualist strategy, but according to which Premise 1 is true at all standards and all the "action" occurs at Premise 2. Mark Heller defends what looks to be such a response in his 1999.)

In addition to conceding to the skeptic that we don't know that we have hands according to the absolute standards, as any response that follows the basic contextualist strategy does, our Moorean contextualist response also concedes that, according to those "absolute" standards, we don't know that we're not BIVs. That's why it cannot just be straightforwardly classified as a response according to which we do know that we're not BIVs. But it is a feature of such contextualist strategies that we do

have "regular," if not "super-high octane" ("absolute"), knowledge of the fact that we're not BIVs. It's in this way that we're Mooreans.

We contextualist Mooreans, then, are not straightforward Moorean antiskeptics, both because we're not straightforwardly Mooreans and because we're not straightforwardly antiskeptics. But in the way that we are Moorean antiskeptics, we are Moorean antiskeptics who avoid the pitfalls of our more straightforwardly Moorean cousins: the Putnam-style responders to skepticism. For ours is a nonheroic response to skepticism and one according to which our knowledge that we're not BIVs is very solid, not highly shaky, knowledge. I'll explain these two features in turn, but only with respect to my own contextualist Moorean account.

First, my response to the skeptical argument is nonheroic. Though I argue that we do know that we're not BIVs (at least according to nonabsolute standards), I don't make this knowledge that I claim we have dependent on any argument I give—or any argument whatsoever. If my argument is right, then not only do I know that I'm not a BIV, but also my Mom knows this about herself—or at least can easily (and without the aid of any fancy argument) come to know this if she were to come to consider the issue of whether she's a BIV.²⁶ As a Moorean, I face the pointed question: *How* do we know that we're not BIVs? I join my contextualist Moorean cohorts in answering negatively: *not* by means of any argument, nor even by means of our having any effective evidence to that effect.²⁷ More positively, in my own case, my account is that we know, according to even quite high standards (though not according to the absolute standards) that we're not BIVs because our belief as to whether we're BIVs matches the fact of the matter in the actual world and in the sufficiently nearby worlds. Gathering evidence and constructing arguments is indeed one way, *but only one way among others*, to improve one's epistemic position with respect to some propositions, and to thereby come to be in a strong epistemic position with respect to them. On my account, we're all—(well, I trust everybody here: we can dream up a character in weird circumstances who is in a poor epistemic position with respect to her not being a BIV)—well-positioned with respect to our not being BIVs before we have any arguments to that effect, and independently of any such arguments we might have.

How well-positioned? This brings us to the second feature: Very well-positioned. Well-enough positioned to not only meet ordinary standards for knowledge, but to be as highly warranted as we take our knowledge of such things as that we have hands to be—though not, of course, well-enough positioned to meet the skeptic's absolute standards. I'm quite explicit about our being in as strong an epistemic position with respect

to "I am not a BIV" as we are with respect to "I have hands" (see section 10 of my 1995). Of course, this somewhat surprising comparative fact can be due either to our being in a surprisingly strong epistemic position with respect to the former or to our being in a surprisingly weak position with respect to the latter. The surprise, on my account, is how well-positioned we are with respect to our not being BIVs. This verdict is defended by means of an account of why it can *seem* that our knowledge that we're not BIVs is shaky at best.²⁸

On my account, there is a conversational mechanism which has the effect that the standards according to which we don't know that we're not BIVs—the "absolute" standards—are the very standards that tend to get put into play by the bringing up of the BIV hypothesis. It's the operation of this mechanism, not actual shakiness, which explains why our knowledge that we're not BIVs can *appear* to be shaky at best.

7. Comparing the Two Moorean Responses to Skepticism

As I warned, my purpose here is not to give a full defense of my alternative response to skepticism. (For a much *fuller*, though still not *full*, general defense of my response, see my 1995.) My purpose has rather been the more limited one of showing how my alternative has potential to fare better with respect to the problems that the externalist solution faces in connection with its heroic nature: There is a nonheroic alternative to the externalist response, even within the Moorean camp, broadly enough construed.²⁹

In brief summary, the two views respond very differently to the apparent shakiness of whatever knowledge we might have to the effect that we're not BIVs, the externalist heroically seeking to shore up this shaky belief, while the contextualist Moorean seeks to argue that the belief is not, and has not been, as shaky as it appears: We really do have quite highly warranted knowledge (even if not absolute knowledge) of the fact that we're not BIVs; that's our natural state. Or, since we may fail to know that we're not BIVs in our "natural state" because we never consider whether we're BIVs and so don't even believe that we're not BIVs,³⁰ our natural state may perhaps more accurately be described as follows: We're naturally disposed to reject such hypotheses as that we're BIVs, and to thereby come to believe *and to know* (*with a high degree of warrant*) that these hypotheses are false, upon coming to consider them.

But whatever one thinks of the prospects of my alternative response, I hope it's clear that the externalist's is quite hopeless as a response to the problem of skepticism.

8. Externalist Arguments Freed from the Burden of Answering Skepticism

As I wrote at the end of section 1, there are two different ways to approach the issue of whether and how we can know that we're not BIVs: One might be interested in such answers in order to see whether they might be used to cut off BIV-inspired skepticism via a denial of the first premise of such skeptical argument, or one might simply be interested in such questions as whether we can have a priori knowledge of such things as that we're not BIVs, regardless of whether such a priori knowledge can be used to combat skeptical arguments. Here we use "a priori" in a broad sense that has become quite popular in discussions of whether a priori knowledge of one's own thoughts is compatible with Semantic Externalism: It is construed so that it includes introspective knowledge of one's own mental states, so long as this introspective knowledge is independent of warrant obtained by empirical investigation of the external world. It's important to stress this because many would use "a priori" in a different way in which introspective knowledge, even if it didn't depend on empirical investigation of the external world, would still not count as a priori because it does depend on "internal," if not on "external," experience. In our broad sense of "a priori," both premises of the Compatibilist argument seem to be at least arguably a priori; that argument, recall, is:

Compatibilist Argument

1. I think that water is a liquid.
 2. No BIV can think that water is a liquid.
- So,
C. I am not a BIV.

Premise 2 is at least arguably a priori because it seems to be supported by thought-experiments of the type Putnam made famous, rather than by any empirical investigation of the world. Premise 1 is at least arguably a priori in our broad sense because it at least seems to be something one can know by introspection, without having to investigate the world. One might be interested in whether one can, by means of such an argument, come to have a priori knowledge of such things as that one is not a BIV. Once she gives up the dream of responding to skepticism by means of such an argument, many of the heaviest burdens of the Compatibilist arguer are lifted.

First, when the Compatibilist arguer is trying to answer BIV-inspired skepticism, it does no good to shoot down just a

few ways of construing the BIV hypothesis. To stop the skeptic, all versions of the hypothesis must be shot down, since the skeptic only has to win once. But when she is simply interested in discovering what we can come to have (broad) a priori knowledge of, it is significant that this or that version of the hypothesis can be known a priori to be false, even if other versions can't be known a priori to be false.

Second, when she's trying to answer the skeptic, the Compatibilist arguer, for reasons we've seen, as a Moorean, has to worry about whether the knowledge gained by means of the argument is too shaky to do the work it has to do in protecting our knowledge that we have hands. By contrast, in simply investigating whether she can attain knowledge that she's not a BIV that is a priori in that it doesn't depend on external observation, with no eye toward answering the skeptic, there's no pressing reason why the Compatibilist arguer would need to limit her interest to whether she can attain very highly warranted knowledge. Indeed, I don't see any reason why her interest should be limited to warrant sufficient for knowledge: She can be interested in whether a priori justification or warrant can be had for the conclusion that she is not a BIV, even if this warrant is not sufficient for knowledge. This is important because one might doubt that whatever knowledge we might get for the likes of Premise 2 by means of thought experiments can really be so highly warranted. I'm a friend of externalism myself, but when I compare the likes of 2 with the likes of "I have hands," the former starts to look awfully shaky to me.

So perhaps we can get a priori knowledge of, or at least a priori warrant for, the claim that we're not BIVs by means of such an argument (though there are many potential obstacles to this that I won't investigate here). However, I suspect that the best ways of filling out and then evaluating my alternative, contextualist Moorean account of how we know that we're not BIVs will have it come out also as an account according to which our knowledge that we're not BIVs is a priori. If so, and if that account is correct, then the Compatibilist argument would not be providing us with any new a priori knowledge of its conclusion: That conclusion would be something we already knew, already knew a priori, and already knew with a very high degree of warrant, long before we ever encountered this relatively shaky argument for it. Still, I see no bar in principle to compatibilist arguments being sound a priori arguments for their conclusion. In that respect, the externalist strategy may have some merit.

9. Problematic A Priori Knowledge

There has been a great deal of interest in recent years in the issue of what a priori knowledge one might be able to gain by

means of Compatibilist arguments closely related to our above Compatibilist argument, but with significantly different conclusions. Like our above argument, these arguments proceed from a premise about the content of one's own thought, together with an externalist premise, but these arguments reach more problematic conclusions about one's environment—conclusions that seem to most people to be such that it's crazy to suppose we can know them a priori. Indeed, these arguments are typically put forward as reductions of Compatibilism, where Compatibilism is the conjunction of the claims (a) that we typically do know our own thought contents a priori (in the broad sense under discussion) with (b) that we know, by means of Putnam- and Burge-like thought experiments, that semantic externalism is correct. The objection to Compatibilism is that if it were correct, we would be able to know a priori facts about our environment that clearly can't be known a priori.³¹ For instance, Paul Boghossian has argued that the Compatibilist is committed to accepting that a person (named Oscar) who is "chemically indifferent" (he doesn't know the chemical composition of water) but who knows that his concept of water does at least aim to pick out a natural kind, has available to him this Compatibilist argument:

Oscar's Compatibilist Argument

1. I have the concept *water*.
2. If I have the concept *water*, then *water* exists.

So,

C. *Water* exists.³²

Boghossian concludes that "the Compatibilist is committed to the manifestly absurd conclusion that we can know a priori that *water* exists."³³ Boghossian is far from alone in his judgment of absurdity: Both attackers and defenders of Compatibilism agree that it's absurd to suppose that you can know the likes of "Water exists" a priori. I don't know if there would be similar agreement about the absurdity of knowing a priori that one is not a BIV. But in any case, in line with my above discussion of potential merit and the limitations of our Compatibilist argument to the conclusion that one is not a BIV, I'd like to close by suggesting a related approach to Compatibilist arguments like Oscar's.

Though I can't argue it here, I think that we have to get to the idea that we can have a priori knowledge of—or at least a priori warrant for—conclusions of Compatibilist arguments that appear to most to be very problematic as potential items of a priori knowledge. The standard line of defense from the *reductio*s that Compatibilists have attempted is to agree that it would be absurd to suppose we have a priori knowledge of those conclusions, but to argue that, when you look carefully

at the externalist premise of the Compatibilist argument, it is not something that is known a priori on plausible accounts of externalism, because, although the premise issues from a thought-experiment rather than directly from empirical investigation, knowledge we have gained through experience—like that water is a natural kind—guides the verdict we make about the thought experiment.³⁴ I believe that when we get the details of externalism straightened out, together with what aspects of it are knowable a priori, it will turn out that the Compatibilist is committed to the a priori knowability, not of anything as simple as “Water exists,” but of other facts about one’s environment, that are more complicated than “Water exists,” but are otherwise no more palatable as objects of a priori knowledge. Perhaps something like:

Either I have causally interacted with water, or water actually is not (as I suppose) a natural kind,

or perhaps (and more realistically, given my sense of what an acceptable form of externalism will end up looking like) something more complicated, but still not much more palatable.

If I’m right about that, then Compatibilists need a different defense. I suggest that we Compatibilists accept that the Compatibilist arguments can give one a priori knowledge of—or at least a priori warrant for—problematic conclusions, but seek to show how this might not be so bad. Several mitigating factors present themselves. First, the conclusions of *these* compatibilist arguments (as opposed to the conclusion that one is not a BIV) are things we clearly know anyway through external experience. Second, we know these conclusions very solidly—with a high degree of warrant—through external experience. Third, whatever warrant we might get of the externalist premise of the compatibilist argument, and thereby get of the conclusion of the argument, through a philosophical thought-experiment is going to be very shaky as compared with our wonderfully solid a posteriori knowledge of that same conclusion, even if it is fairly solid *as far as philosophical arguments go*. In light of the fact that we already know the conclusions of these Compatibilist arguments with a very high degree of warrant through empirical means, is it really so absurd to suppose we might later also come to have some relatively very minor and very shaky a priori warrant for them? It doesn’t seem that absurd to me. Which is good, because, as I’ve already noted, I think that, seemingly absurd or not, this is something that we’re going to have to learn to live with.³⁵

Notes

¹ Indeed, in my own experience of questioning introductory philosophy students, I think a majority are inclined to think they do

know that they’re not BIVs. I must point out, however, that that’s my own dominant intuition about the question, and also *in a way*—see section 6—my settled opinion on the matter. So, though I try to remain neutral as I elicit “intuitions,” I might well be, despite my best efforts, presenting the question in a way that unfairly favors the answer I prefer, and this question, even more than other philosophy questions one asks about, seems to be one where the way it is posed plays an important role in how it gets answered.

See Sosa (1999), which reports quite similar results from “informal polling” (note 16). Sosa registers an impression very close to my own regarding how intuitively unclear it is whether we know that we’re not BIVs (1999, 147). Though my own dominant inclination, as reported above, initially was to say that I do know I’m not a BIV, for a period of time I was so immersed in the task of explaining the plausibility of the skeptic’s claim to the effect that we don’t know this that I forgot, as Sosa puts it, how simultaneously *implausible* that same skeptical claim is. It was a comment from John Greco that served as the needed slap in the face to remind me how unclear things are—and even that my own intuition goes against the skeptic!

² I call these “Moorean” solutions because G. E. Moore famously responded to the *dream argument*—the instance of our Argument by Skeptical Hypothesis where the skeptical hypothesis is the possibility that one is dreaming—by denying its first premise. Whether Moore would have responded in a similar way to the BIV argument, or, to take another ultraradical hypothesis (and one Moore surely would have thought about), Descartes’s evil genius argument, is very far from clear. Moore’s own way of carrying out the “Moorean” response (very briefly: by claiming that one indeed has conclusive evidence that the hypothesis is false, though one cannot say what that conclusive evidence is), while it’s already highly problematic as a response to the dream argument, is worse as applied to the skeptical arguments employing the radical hypotheses.

³ Putnam, 1981, 8/*Skepticism*, 32.

⁴ Of the BIV scenario, Putnam writes early in his essay: “When this sort of possibility is mentioned in a lecture on the Theory of Knowledge, the purpose, of course, is to raise the classical problem of scepticism with respect to the external world in a modern way. (How do you know you aren’t in this predicament?) But this predicament is also a useful device for raising issues about the mind/world relationship” (1981, 6/*Skepticism*, 31). And in what follows, Putnam’s own interest seems confined to the “mind/world relationship,” for the “classical problem of scepticism” is hardly ever again mentioned.

⁵ Putnam, 1981, 14/*Skepticism*, 36–37

⁶ The first form of argument—or what I’m about to call the “Dilemma” version—seems closer to most of what Putnam wrote. The second, “Compatibilist” form of the argument is at least suggested by passages like: “if we can consider whether it [the supposition that we are brains in a vat] is true or false, then it is not true (I shall show). Hence, it is not true” (1981, 8/*Skepticism*, 32). The most obvious candidate for the implicit premise that Putnam is assuming here, which can combine with his explicit premise, *If we can consider it, then it is not true*, to reach his conclusion, *It is not true*, is, of course, “*We can consider it*. If that’s what he has in mind here, then Putnam is giving a version of the Compatibilist Argument.

⁷ Readers will have noticed a slide in the argument between "I am not a BIV" is true" and "I am not a BIV." The argument is often reconstructed to include a separate premise licensing this move of "disquotation," and much of the controversy about the Dilemma argument is over the legitimacy of the needed disquotation.

⁸ For reasons we'll see in section 9.

⁹ Putnam (1975) and Burge (1979, 1982, and 1986).

¹⁰ Putnam, 1981, 15/*Skepticism*, 7; emphasis added in both quotations.

¹¹ In my 1999a, I wrote that this problem with the antiskeptical strategy was "first noted" (9) in Brueckner (1986). That was a mistake. Smith (1984), for instance, not only notes the general problem that Putnam's strategy works only against some forms of the BIV hypothesis, but also suggests the specific possibility of recent envatment as skeptical hypothesis that would avoid the Putnamian counterattack. I don't know if Smith (1984) was the first place these things were done. My understanding is that these were often noted in the explosion of literature that followed Putnam's work. They were ideas "in the air."

¹² Putnam, 1981, 6/*Skepticism*, 31.

¹³ Christensen, 1993, 302, fn. 1.

¹⁴ Warfield, 1998, 142–143/*Skepticism*, 88.

¹⁵ Well, that's a *bit* too quick. Since certain stretches of our presumed knowledge of the external world may escape being undermined by certain skeptical hypotheses, the skeptic may need several hypotheses to work in order to achieve the scope of the skepticism she desires. Still, it is true that to undermine any given piece or stretch of external world knowledge, the skeptic needs just one skeptical hypothesis to work against the presumed knowledge in question. It's pretty clear that enough skeptical hypotheses will escape Putnam-style refutation to—if there's no other problem with the skeptic's argument—undermine an unnervingly vast scope of our presumed knowledge. As far as your knowledge as it's protected by the Putnam strategy goes, your hands are gone—that's for sure. Perhaps, though, something like *There is an external physical world* will be protected, and that would not be an insignificant antiskeptical result. Still, we want a defense that protects our more detailed knowledge of the external world—like that we have hands—don't we?

¹⁶ The Putnam-style answer to skepticism is a current manifestation of a tradition of approaches to skepticism that has been historically very prominent—approaches that attempt to *refute* the skeptic, where a refutation requires an argument that derives an antiskeptical conclusion from premises that in an important way *don't beg the question* against the skeptic. I belong to a completely different tradition—one which, in a sense I cannot here take the space to explain, attempts to *defeat* the skeptic, without refuting her. Because of our antiskeptical campaigns that are off-limits to would-be refuters of skepticism, our approach can appear to be much weaker. Showing how the Putnam-style approach would provide only a weak response to skepticism *even if it did work* is especially important to me because it illustrates a very *general* weakness in attempted refutations of skepticism. In some very important ways, refutations of skepticism, even if they did work (which none do), would be of very limited value.

¹⁷ Here, I'm assuming that the Argument by Skeptical Hypothesis can be applied to the nonphilosophical. And it can be. But the skeptic will probably want to adjust the argument a bit: She can't just replace "I" with "My Mom" throughout to obtain:

1'. My Mom doesn't know that not-H.

2'. If my Mom doesn't know that not-H, then she doesn't know that O.

So, C'. My Mom doesn't know that O.

The problem is that my Mom may fail to even believe that not-H because she has never even considered the issue of whether H is true. When we keep in mind that way of failing to know that not-H, then the second premise loses plausibility. To regain that lost plausibility, the skeptic will have to modify the second premise to something like:

2". If my Mom doesn't know that not-H, and if (should she come to consider whether H) she couldn't get to know that not-H by deducing that not-H from her belief that O, then she doesn't know that O.

Having weakened the second premise in that way, the skeptic will then have to compensate by strengthening the first premise to something like:

1". My Mom doesn't know that not-H, and (should she come to consider whether H) she couldn't get to know that not-H by deducing that not-H from her belief that O.

It may look as if the skeptic loses plausibility by strengthening 1' to 1". In particular, it can appear that, in so doing, the skeptic begins to beg an important question against the Moorean, by beginning to assume that we cannot get to know that not-H by deducing that not-H from O. But in the way the skeptic begs the question by her use of 1", she has been begging the question all along against Mooreans. The Moorean denies 1", and even has a valid argument, from intuitively plausible premises, against 1", namely: not-C'; 2"; therefore, not-1". Against this, the skeptic puts forward the intuitive plausibility of 1". If this is begging the question, the skeptic already begged the question against the Moorean in our *original* formulation of the Argument. In using premise 1, the skeptic begs the question against the Moorean, who denies 1, and who even presents this argument, from intuitively plausible premises, against 1: not-C; 2; therefore, not-1. (Of course, in using not-C as a premise against the skeptic, the Moorean in some sense "begs the question" against the skeptic. I guess the fair verdict to reach is that in this way, both parties beg a question against the other.) Indeed, already in our original formulation, where the argument was applied to me, rather than my Mom, the plausibility of 1 already depended on our assuming that I couldn't get to know that not-H by deducing that not-H from O. The only reason that we were able to get away with keeping this as an implicit assumption and thereby keeping the argument simple when it was applied to me was that we all were aware of the fact that I was aware of the BIV hypothesis, and so, if it

were possible, would have come to know that not-H by deducing it from O. So, all along, the skeptic has been assuming that folks can't come to know that not-H simply by deducing that not-H from their belief that O, and so has been "begging" that question. Still, I find this assumption, and 1 and 1", to enjoy a good deal of intuitive plausibility, though, as I've already admitted in the text, I don't find all this as intuitively plausible as some philosophers seem to. (Thanks to Michael DePaul for pointing out to me the need to address this issue in this paper.)

¹⁸ Anthony Brueckner's treatment of the strategy is fairly typical. Brueckner presents the strategy as a response to "Cartesian skepticism," which consists of a skeptical argument much like the one we're dealing with. In particular, the skeptical argument as he formulates it contains the premise "I do not know that I am not a BIV" (1992b, 200/*Skepticism*, 43)—our premise 1, where the skeptical hypothesis is the BIV scenario—and it's this premise, which he says is targeted by the semantic antiskeptical strategy, which is described simply as "blocking the skeptical argument" (202/*Skepticism*, 44), apparently at the premise we've identified. He makes some of the preliminary moves of the antiskeptical argument, and then asks the crucial question: "How does any of this help with the refutation of Cartesian skepticism?" (204/*Skepticism*, 46), but then just dives into the details of the argument, without explaining how such an antiskeptical argument would help in disarming the skeptical argument. After presenting the argument he wishes to discuss, which has as its conclusion, "I am not a BIV" (205/*Skepticism*, 48), he writes, "So if we accept the externalist semantic assumptions underlying the argument, it seems that we *do* get the desired antiskeptical conclusion that I am not a BIV" (205/*Skepticism*, 48). He then goes on, in the rest of his paper, to evaluate whether we really do, as it seems, get this conclusion, but doesn't say how getting this conclusion "blocks" the skeptical argument. Maybe it's supposed to be obvious. But the only *obvious* way I can see that "getting" this conclusion derails the skeptical argument is by, if need be, giving us knowledge that we're not BIVs and thereby making the premise under discussion false if it isn't already false. There is no obvious way that "getting" this conclusion shows that I knew all along that I was not a BIV, or that my Mom, who has not encountered the argument and so has not "gotten" the conclusion, knows that she's not a BIV.

¹⁹ I have some ideas about how this might best be done—how one might try to use Putnam-style externalist considerations in a nonheroic response. I don't think that such a strategy will work in the end, but I should leave its formulation to those more sympathetic to the approach.

Timothy Williamson, in his comments which follow, presents a very interesting response of this type, and argues that aspects of my own views license such a response. Replying to this suggestion will have to await another occasion.

Forbes (1995) is an interesting case. Forbes presents his Putnam-style argument in a heroic mode; he writes, for instance: "If we have followed Putnam's argument in a competent way, we *know* the conclusion ['I am not a BIV'] is true" (207/*Skepticism*, 63). But he follows a strategy, which, by arguing that there are no relevant alternative situations in which one's belief that one is not a BIV is

false, seems to hold promise for securing the knowledge of the unphilosophical who have not followed Putnam's argument in any way. I will not here evaluate the prospects of this as a nonheroic response except to make two points. First, for reasons I cannot go into here, I don't think Forbes's strategy will work without a fairly (and perhaps objectionably) strong form of Semantic Externalism—perhaps even what we're calling "high-grade externalism"—underwriting it. It certainly won't work (at least without serious modifications I can't see how to make) if Putnam's suspicion (which I take seriously) that the BIVs fail to have any truth-evaluable thoughts at all should be right. On this, see Forbes's admission in fn. 3 (1995, 207/*Skepticism*, 73–74). Second, and more important for our purposes, it is clear at any rate that this strategy won't work against many of the versions of the BIV hypothesis—like recent envatment. Forbes would not deny this, and does not even intend his strategy to target the more difficult formulations of the hypothesis (see 1995, 207/*Skepticism*, 63).

²⁰ Here in the text I suppress my own contextualist views for the time being for ease of exposition. As a contextualist, I believe that the strength of the call of "knowledge" varies with context, and that there are contexts in which, strongly warranted as it is, my belief that I have hands does not count as knowledge. What I mean, taking my contextualism into account, is that my knowledge that I have hands is warranted to an extent that far exceeds what's needed for knowledge in most ordinary contexts.

²¹ Of course, if things don't work out, one may take the lesson of our investigation to be that we shouldn't have set that option aside!

²² See Stine (1976), Cohen (1987, and especially 1988), and Lewis (1996).

²³ Explained at somewhat greater length in section 2 of my 1995 (4–7/*Skepticism*, 185–187).

²⁴ For an explanation of such theories, see my 1992 and my 1999b.

²⁵ The skeptic we resist, then, is what I call the "bold" skeptic (see 1995, 5–6/*Skepticism*, 185–186).

²⁶ See notes 17 and 30. I imagine this is how my Mom would come to believe (and, according to me, know) that she's not a BIV if the issue were pressed on her: Immediately upon hearing the hypothesis that she's a BIV, she would reject it as silly.

²⁷ Thus, writing about the skeptical hypothesis that the "zebra" she sees is really a cleverly painted mule, though I suppose she would take the same line on the BIV hypothesis, Stine writes:

[O]ne does know what one takes for granted in normal circumstances. I do know that it is not a mule painted to look like a zebra. I do not need evidence for such a proposition. The evidence picture of knowledge has been carried too far.... [I]f the negation of a proposition is not a relevant alternative, then I know it—obviously, without needing to provide evidence.... (1976, 258/*Skepticism*, 153)

Along similar lines, Lewis writes:

Do I claim you can know *P* just by presupposing it?! Do I claim you can know that a possibility *W* does not obtain just by

ignoring it? Is that not what my analysis implies, provided that the presupposing and the ignoring are proper? Wall, yes. (1996, 561–562/*Skepticism*, 233)

Here, “I am not a BIV” is certainly an example of a *P* which, on Lewis’s view, one can know “just by [properly] presupposing it,” and “I am a BIV” is an example of a *W* which, on his view, you can know is false “just by [properly] ignoring it.” Now, Lewis has a lot to say about what can be *properly* presupposed and *properly* ignored, and we can’t get into all that here. For our purposes, suffice it to say that arguments and/or evidence are not needed for us to properly ignore the possibility that we’re BIVs, to properly presuppose that we’re not BIVs, and to in that way know that we’re not BIVs.

Similarly, for Cohen, our beliefs in the nonobtaining of radical skeptical hypotheses, like the BIV hypothesis, are “intrinsically rational,” by which Cohen means a belief that one rationally holds “without possessing evidence for the belief” (1988, 112). Cohen makes it clear (see especially 113–115) that our beliefs that we’re not BIVs are rational enough to count as knowledge (at least at ordinarily moderate standards for knowledge), and that we do indeed know, according to moderate standards, that we’re not BIVs, without our having evidence to that effect. (Note that Cohen offers quite a different account of how we know that nonradical, or moderate, skeptical hypotheses are false. But he explicitly counts the BIV hypothesis as a “radical” one, so we can be sure that it is a hypothesis for which Cohen would give an “intrinsically rational” account of how we know it’s false.)

²⁸ Actually, in my 1995, I’m trying to account for why it can seem that we don’t know at all that we’re not BIVs. But that account would work as well to explain the somewhat “broader” phenomenon I’m here seeking to explain: why it can seem that our “knowledge” that we’re not BIVs is shaky at best—is either not knowledge at all, or, if it is knowledge, is just barely knowledge.

²⁹ Sosa (1999) presents another, quite different, Moorean response, which shares with the contextualist Moorean response the virtue of being nonheroic.

³⁰ See note 17. It’s not completely clear to me that those who have never considered the BIV hypothesis don’t know that it’s false. Perhaps they have “implicit” knowledge of the falsehood of the hypothesis.

³¹ See McKinsey (1991), Brown (1995), and Boghossian (1997).

³² Boghossian, 1997, 165.

³³ Boghossian, 1997, 166. These quotations are a bit misleading in one respect: Boghossian recognizes that it is not the simple “Water exists,” but rather something a bit more complicated, but, according to Boghossian, no more palatable, than the externalist will have to admit as a priori knowable; see especially 166–167.

³⁴ See Brueckner (1992a), McLaughlin and Tye (1998a), and Part 2 of McLaughlin and Tye (1998b).

³⁵ Thanks to Timothy Williamson for enlightening discussion about these matters.

References

- Boghossian, Paul. 1997. What the externalist can know a priori. *Proceedings of the Aristotelian Society* 97:161–175.
- Brown, Jessica. 1995. The incompatibility of anti-individualism and privileged access. *Analysis* 55:149–156.
- Brueckner, Anthony. 1986. Brains in a vat. *Journal of Philosophy* 84:148–167.
- . 1992a. What an anti-individualist knows a priori. *Analysis* 52:111–118.
- . 1992b. Semantic answers to skepticism. *Pacific Philosophical Quarterly* 73:200–219. Reprinted in *Skepticism*, edited by K. DeRose and T. Warfield, Oxford University Press, 1999.
- Burge, Tyler. 1979. Individualism and the mental. *Midwest Studies in Philosophy* 4:73–121.
- . 1982. Other bodies. In *Thought and Object: Essays on Intentionality*, edited by A. Woodfield. Oxford University Press.
- . 1986. Individualism and Psychology. *Philosophical Review* 95:3–45.
- Christensen, David. 1993. Skeptical problems, semantical solutions. *Philosophy and Phenomenological Research* 53:301–321.
- Cohen, Stewart. 1987. Knowledge, context, and social standards. *Synthese* 73:3–26.
- . 1988. How to be a fallibilist. *Philosophical Perspectives* 2:91–123.
- DeRose, Keith. 1992. Contextualism and knowledge attributions. *Philosophy and Phenomenological Research* 52:913–929.
- . 1995. Solving the skeptical problem. *Philosophical Review* 104:1–52. Reprinted in *Skepticism*, edited by K. DeRose and T. Warfield, Oxford University Press, 1999.
- . 1999a. Responding to skepticism. In *Skepticism*, edited by K. DeRose and T. Warfield, Oxford University Press, 1999, 1–24.
- . 1999b. Contextualism: An explanation and defense. In *The Blackwell Guide to Epistemology*, edited by J. Greco and E. Sosa. CITY: Blackwell Publishers, 187–205.
- DeRose, Keith, and Ted Warfield, eds. 1999. *Skepticism: A contemporary reader*. Oxford University Press.
- Forbes, Graeme. 1995. Realism and skepticism: Brains in a vat revisited. *Journal of Philosophy* 92:205–222. Reprinted in *Skepticism*, edited by K. DeRose and T. Warfield, Oxford University Press, 1999.
- Heller, Mark. 1999. Relevant alternatives and closure. *Australasian Journal of Philosophy* 77:96–208.
- Lewis, David. 1996. Elusive Knowledge. *Australasian Journal of Philosophy* 74:549–567. Reprinted in *Skepticism*, edited by K. DeRose and T. Warfield, Oxford University Press, 1999.
- Ludlow, Peter, and Noah Martin. 1998. *Externalism and self-knowledge*. CSLI Publications.
- McKinsey, Michael. 1991. Anti-individualism and privileged access. *Analysis* 51:9–16. Reprinted in *Externalism and self-knowledge*, edited by P. Ludlow and M. Martin, 1998.
- McLaughlin, Brian, and Michael Tye. 1998a. The Brown-McKinsey charge of inconsistency. In *Knowing Our Own Minds: Essays in*

- Self-Knowledge*, edited by McDonald, Smith, and Wright. Oxford University Press. Reprinted in *Externalism and self-knowledge*, edited by P. Ludlow and N. Martin, 1998.
- McLaughlin, Brian, and Michael Tye. 1998b. Is content externalism compatible with privileged access? *Philosophical Review* 107:349–380.
- Putnam, Hilary. 1975. The meaning of meaning. In H. Putnam, *Mind, Language, and Reality: Philosophical Papers, Volume II*. Cambridge University Press.
- . 1981. Brains in a vat. *Reason, Truth, and History*. Cambridge University Press. Reprinted in *Skepticism*, edited by K. DeRose and T. Warfield, Oxford University Press, 1999.
- Smith, Peter. 1984. Could we be brains in a vat? *Canadian Journal of Philosophy* 14:115–123.
- Sosa, Ernest. 1999. How to defeat opposition to Moore. *Philosophical Perspectives* 13:141–153.
- Stine, Gail. 1976. Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies* 29:249–261. Reprinted in *Skepticism*, edited by K. DeRose and T. Warfield, Oxford University Press, 1999.
- Warfield, Ted A. 1998. A priori knowledge of the world: Knowing the world by knowing our minds. *Philosophical Studies* 92:127–147. Reprinted in *Skepticism*, edited by K. DeRose and T. Warfield, Oxford University Press, 1999.

Skepticism, Semantic Externalism, and Keith's Mom

Timothy Williamson
University of Edinburgh

A canonical argument for skepticism about our ordinary knowledge of the external world is this: we have ordinary knowledge of the external world only if we know that we are not brains-in-vats (BIVs); we do not know that we are not BIVs; therefore, we do not have ordinary knowledge of the external world. Hilary Putnam (1981) famously suggested that semantic externalism undermines the skeptical argument. The semantic externalist postulates necessary connections between the content of our thinking and the external world. According to Putnam, those necessary connections preclude BIVs from entertaining the proposition that they are BIVs. If so, there is a valid argument from semantic externalism and the plausible premise that we do entertain the proposition that we are BIVs to the conclusion that we are not BIVs.

One standard objection to Putnam's argument is that for some of the less radical BIV scenarios, no plausible form of semantic externalism precludes BIVs from entertaining the proposition that they are BIVs. Perhaps you lack the requisite connections with the external world if you have always been a BIV, but surely you still have them if you were envatted only last night. Thus, as Keith DeRose notes in his illuminating paper, the semantic externalist argument against skepticism looks insufficiently general. DeRose also raises a new problem for the semantic externalist strategy. That is the main topic of this response.

Consider senses of "semantic externalism" and "BIV" in which semantic externalism does entail that BIVs cannot

Tim Williamson is Professor Logic and Metaphysics at the University of Edinburgh, and has just been elected to the Wykeham Chair of Logic at the University of Oxford. He is the author of Identity and Discrimination (1990), Vagueness (1994), Knowledge and Its Limits (2000) and articles in Mind, Journal of Philosophy, Philosophical Review, Philosophy and Phenomenological Research, Journal of Symbolic Logic, Journal of Philosophical Logic, and other journals. His current interests include knowledge-centered epistemology and the metaphysics of quantified modal logic.