

CHAPTER 3:

TWO SUBSTANTIVELY MOOREAN RESPONSES AND THE PROJECT OF REFUTING SKEPTICISM

CONTENTS

1. Substantively Moorean Responses to AI, Straightforward and Contextualist, and Our Pointed Question: How Can We Know that We're not Brains in Vats?.....	1
2. The Project of Refuting Skeptics—and Anti-Skeptics who Are Not Engaged in that Project.....	3
3. Putnam-Style Responses to AI from Semantic Externalism	8
4. Two Forms of Arguments from Semantic Externalism	13
5. Old Objections: Varieties of Semantic Externalism and Varieties of Skeptical Hypotheses.....	14
6. The Disadvantages of Heroism	19
7. The Challenges Facing Non-Heroic, Moorean Alternative Responses.....	26
8. Comparing the Two Moorean Responses to Skepticism.....	28

TWO SUBSTANTIVELY MOOREAN RESPONSES AND THE PROJECT OF REFUTING SKEPTICISM

1. SUBSTANTIVELY MOOREAN RESPONSES TO AI, STRAIGHTFORWARD AND CONTEXTUALIST, AND OUR POINTED QUESTION: HOW CAN WE KNOW THAT WE'RE NOT BRAINS IN VATS?

Responses to AI (the skeptical argument we wrestled with in SSP) can be classified into broad camps according to whether they deny AI's first premise, deny its second premise, or accept its skeptical conclusion. As I mentioned in section 2 of the previous chapter, because G.E. Moore prominently walked this path, responses that deny AI's first premise can be called 'Moorean' responses, in a substantive, as opposed to methodological, sense. Relative to the skeptical hypothesis that one is a BIV, then, the substantively Moorean response involves the claim that one does indeed know that one is not a BIV.

But this way of classifying responses applies most smoothly only to what I have called 'straightforward,' i.e., non-contextualist, responses. As we saw in SSP, the verdicts contextualists give to the questions key to this way of classifying responses (is the first premise, the second premise, the conclusion, true?) can be complicated by the contextualist's varying epistemic standards. Still, the response of SSP can be profitably classified as a *contextualist Moorean* response to AI, and in this it is like the responses of Stewart Cohen (1987; 1988), and David Lewis (1996)—and also Gail Stine (1976), whose pioneering work we are all following here.

First, ours are contextualist responses: We follow the basic contextualist strategy, described in section 2 of SSP: We accept a contextualist theory of knowledge attributions; we accept that at some unusually high standards for

knowledge, 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 AI, at least in large part, by claiming that the presentation of the skeptic's argument has at least some tendency to put into play the very 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 standards that the presentation of AI tends to put in play, ours is not a thoroughly anti-skeptical response. The key way in which ours is an anti-skeptical response is that we do seek to protect the claim that we do know such things according to many non-absolute standards, and thus seek to protect the truth values of our ordinary claims to know.

Second, then, in the way that ours is an anti-skeptical response at all, it is a Moorean response, because, according to us, at the ordinary standards at which we know we have hands, we also know that we are not BIVs. Where the occurrences of 'know' in AI are read in a consistent way that makes the argument unsound, it's at its first premise that it fails. 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 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 Dretskean/Nozickean' response to AI: a response which 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 such a response in his (1999).)

In addition to conceding to the skeptic that we don't know that we have hands according to some unusually high standards, as any response that follows the basic contextualist strategy does, our Moorean contextualist response also concedes that, according to those same high 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 simply do know that we're not BIVs. But it is a Moorean feature of our contextualist strategies that, according to us, we do know by ordinary standards that we're not BIVs. And, as I mentioned in section 9 of the previous chapter, it's a feature of my own view that we know that we're not BIVs by considerably higher-than-ordinary standards, too—though not, of course, by the extraordinarily high standards for knowledge that the presentation of the skeptical argument has at least some tendency to put into play.

Those most sensitive to the appeal of skepticism often raise the question in the title of this section against Moorean responses in a pointed way: 'How can we *possibly* know that we're not BIVs?' The claim of contextualist Mooreans that we do have 'regular,' even if not 'super-high octane,' knowledge of the fact that we're not BIVs exposes us to a version of that question, since we can be asked, perhaps pointedly: How can we possibly know, *even by just ordinary standards for knowledge*, that we're not BIVs? At the same time, it exposes us to rebuke from our fellow Mooreans who desire a more robustly anti-skeptical stance according to which we *simply know* that we're not BIVs—and simply know that we have hands.¹ Such is the lot of contextualists. I answer this pointed question in Chapter 7 of this book. In this chapter, we examine substantive Mooreans who answer it in a very different way.

2. THE PROJECT OF REFUTING SKEPTICS—AND ANTI-SKEPTICS WHO ARE NOT ENGAGED IN THAT PROJECT

Contextualist Moorean responses to skepticism like that attempted in SSP can strike many as feeble. First, of course, there's the whole contextualist aspect of it that we just looked at: conceding to the skeptic that we don't know that not-H and don't

¹ Here I use 'simply know' in the way I did in (DeRose 2009: e.g., 70): To *simply know* something in this semi-technical sense is to be such that any speaker of ordinary English, whatever their context, would speak the truth in saying that you 'know' the thing in question, and would be saying something false if she denied that you 'know.' This could be either because invariantism about 'know(s)' is correct and one meets the standards that invariably govern knowledge attributions, or because, though contextualism is true, one meets all the allowable standards.

know that O according to the high epistemic standards that the skeptic is using. These contextualist aspects of our solution will be discussed in the following chapter.

But then there's also the whole conservative, methodologically Moorean, 'defeating' approach to dealing with the skeptic that I explained and defended in the previous chapter. My defense of this approach was direct, achieved by citing the connection that 'defeating' skepticism in a Moorean fashion has with the philosophically important goal of rationally governing one's acceptances in response to reasons and arguments. Still, I realize that to many the attempt to merely 'defeat' the skeptic in this Moorean way can seem anemic. It can feel too much like trying to eek out a narrow victory against the skeptic on points. Don't we need a more conclusive smack-down of the skeptic? In particular, wouldn't an acceptably robust response involve *refuting* the skeptic in a way that those who follow Moore's methodological approach don't even attempt?

In this chapter, then, we will engage in some comparison, evaluating the defeating/contextualist/methodologically-Moorean approach of SSP against another approach within the substantively Moorean camp, prominent in fairly recent decades, that makes no use of contextualism and that does attempt the kind of refutation of skepticism in question. While such a refutation might initially seem to provide a more satisfying and robust answer to skepticism, I will point out important ways in which such a refutation, even if it could be provided, would in important ways be a weak answer to skepticism, failing to provide much of what we should want from a good response.

As noted in the previous chapter, to refute a skeptic, on the fairly natural use of 'refute' in question, is to derive an anti-skeptical result from premises that *do not beg the question against the skeptic*. In the case of the AI skeptic, such an anti-skeptical result would be that one knows that O, or, more likely, that O is true, or at least that H is false. What it is to 'beg the question' in an argument is very difficult to make precise. But the basic idea here is that the premises of the anti-skeptical argument should not be matters under dispute in the battle between the skeptic and the anti-skeptic—they must be premises that the skeptic either does or in some fairly strong sense must (since she's committed to them) grant. Since the skeptic in view is often the 'external world skeptic,' such premises are often thought to be limited to certain claims about one's own mental states and perhaps certain

metaphysically necessary truths that one can ‘see’ to be true in an *a priori* way. At any rate, claims about the ‘external world,’ and claims to the effect that one knows something about the external world are taken to be clearly out-of-bounds as starting points for attempted ‘refutations’ of the skeptic.

On something like this natural construal, the project of refuting skepticism has been front-and-center in the history of anti-skepticism. Focusing on Kant, but remarking about a major chunk of early modern wrestling with skepticism, Paul Guyer writes:

Kant clearly conceived of the problem of knowledge in terms of methodological solipsism. That is, like thinkers from Descartes to Hume, he supposed that an answer to skepticism must lie in what reflection on his own thoughts can reveal to the individual thinker even on the supposition that nothing but his own consciousness exists. (Guyer 1979: 161)

And in the surrounding context of the above, Guyer makes it clear that the reason the great dead philosophers thought the anti-skeptic must limit her resources in this way is to avoid ‘question-begging’: to reach beyond these tight limits for argumentative starting points is to rest on disputed matters. And as we’re about to see, attempts to in this way ‘refute’ skepticism have continued to fascinate us in much more recent times.

In fact, some may wonder what else could even qualify as a credible anti-skeptical project. Not only front-and-center, the project of refuting the skeptic can appear, at least to some, to be the only worthwhile anti-skeptical game in town. This appearance goes a long way toward explaining why responses to skepticism that eschew the attempt to refute skepticism—of the contextualist-Moorean variety, but also of other types—can be so exasperating to skeptics and to their friends and defenders.

But, exasperating or not, non-refuting retorts to skepticism have abounded since Moore, and perhaps largely due to (yet another aspect of) his influence. In his most famous anti-skeptical work, ‘Proof of an External World’ (Moore 1959a), Moore plays on historical expectations, opening the lecture/essay by quoting Kant (from the Norman Kemp Smith translation) as writing

It still remains a scandal to philosophy . . . that the existence of things outside us . . . must be accepted merely on *faith*, and that, if anyone thinks good to doubt their existence, we are unable to counter his doubts by any satisfactory proof. (Moore 1959a: 127)

Moore then ostensibly sets out to provide the needed proof. But while the reader has been led to expect a refutation of skepticism, the ‘proof’ produced proceeds from premises—‘Here is one hand ... and here is another’ (1959a: 146)—that are just the sort of starting points that must be avoided if one is attempting such a ‘refutation.’ Moore defends his argument as being a ‘perfectly rigorous proof’ (1959a: 145), but his criteria for a ‘proof’ are just that (1) the premise(s) be different from the conclusion, (2) that the premise(s) be things he knows to be the case, and (3) that the conclusion follow from the premise(s) (1959a: 146). While, as I’ve admitted, it is difficult to give a precise account of what constitutes ‘begging the question,’ Moore’s condition (1) is far too weak to effectively rule out question-begging ‘proofs,’ and the ‘proof’ he himself offers, despite satisfying his first condition, does seem to beg the question against the skeptic in just the way supposed to be forbidden to ‘refutations.’ Indeed, it does so obviously enough that one suspects that that was really the point—that Moore is here *intentionally flouting* the expectations of what an anti-skeptical ‘proof’ will be like.

Perhaps Moore would have annoyingly used ‘refutation’ much as he did ‘proof,’ insisting that he was offering a ‘perfectly rigorous refutation’ of skepticism, too, but those who have followed him tend to admit that we are not offering a ‘refutation,’ but are dealing with the problem of skepticism in some other way. For example, here’s Roderick Chisholm:

What we have been saying, of course, is not likely to convince the skeptics and we can hardly claim to have ‘refuted’ them. But our question was not, ‘Can we refute the skeptics?’ Our question was: ‘Are there positive reasons for being skeptical about the possibility of succeeding in the epistemic enterprise?’ The answer seems to be that there are no such reasons. And therefore it is not unreasonable for us to continue. (Chisholm 1989: 4)

More recently, here’s James Pryor, from the first pages of his paper ‘The Skeptic and the Dogmatist,’ explaining two ways of addressing skepticism, the latter of which is his:

One might go about grappling with such a skeptic in two different ways.

The **ambitious anti-skeptical project** is to refute the skeptic on his own terms, that is, to establish that we can justifiably believe and know such things as that there is a hand, using only premises that the skeptic allows us to use. The prospects for this ambitious anti-skeptical project seem somewhat dim.

The **modest anti-skeptical project** is to establish to our satisfaction that we can justifiably believe and know such things as that there is a hand, without contradicting obvious facts about perception. This is not easy to do, for the skeptic can present us with arguments from premises we find intuitively acceptable to the conclusion that we cannot justifiably believe or know such things. So we have a problem: premises we find plausible seem to support a conclusion we find unacceptable. The modest anti-skeptical project attempts to diagnose and defuse those skeptical arguments; to show how to retain as many of our pretheoretical beliefs about perception as possible, without accepting the premises the skeptic needs for his argument. Since this modest anti-skeptical project just aims to set our own minds at ease, it's not a condition for succeeding at it that we restrict ourselves to only making assumptions that the skeptic would accept. (Pryor 2000: 517-18)

Pryor's modest skeptical project has much in common with my proposed attempt to 'defeat' the skeptic, in the way described in the previous chapter. But most noteworthy to us right now is Pryor's specification of what he is *not* going to be attempting: 'refuting' the skeptic, as we can put it in our use of that term. And, finally, one of Pryor's teachers (and also a philosopher who has greatly influenced my approach), Robert Nozick:

My purpose is not to refute the skeptic, to prove he is wrong, to convince him, to marshal arguments and reasons which must convince him (if he is rational). (Nozick 1981: 15)

Nozick goes on to tightly tie his use of 'refute' to the matter of begging of the question against the skeptic. On the next page, Nozick explains that his own method of engagement with skepticism: by explaining how knowledge is possible in light of the skeptic's argument. The contrast he draws is that in employing his own method, he uses claims that could not be used in the alternative, refuting/proving wrong method, precisely because they would be question-begging: "The goal of explanation makes it legitimate for the philosopher to introduce statements as hypotheses (acceptable to him) that the goal of proof would exclude as begging the question (of proving to the skeptic that he knows)" (1981: 16). I think we should use 'proof' / 'prove' as Nozick does, admitting that we are offering no 'proof' that the skeptic is wrong. And I suspect most working in this broad and diffuse non-refuting tradition of treatments of skepticism are similarly inclined. Rather than insisting that we are 'refuting' the skeptic or 'proving' her wrong, it is more characteristic of this broad tradition (that follows Moore in substance and in spirit here, but not in his use of some key terms) to admit that that's not at all what we're doing, and perhaps even to admit that it can't be done, but to ask, 'So what?'

How great would it be if we did refute the skeptic? Maybe not as great as you suppose, or so I will suggest. At least there would be important anti-skeptical work still to be done even if we had a refutation of skepticism.²

3. PUTNAM-STYLE RESPONSES TO AI FROM SEMANTIC EXTERNALISM

The attempted refutations of skepticism we will consider are Putnam-style attempts to prove that one is not a BIV. For quite some time after Hilary Putnam's 'Brains in a Vat' (in Putnam 1981) got it all going, a prominent answer to our chapter's title question, 'How can we know that we're not BIVs?', was: via a Putnam-style argument from semantic externalism. In calling this answer 'prominent,' I don't mean to suggest it was very widely accepted. In fact, most of the literature on this approach was critical. But while Putnam-inspired attempts to deal with skepticism were not widely viewed as successful, what they at least could be seen as *attempting* to achieve—a *refutation* of BIV-inspired skepticism, as we're using the term here—may have done much to set expectations for what a sufficiently robust anti-skeptical project would look like—though, as we've seen, high expectations for anti-skeptical projects have a long and storied history in philosophy. Contextualist Mooreanism could seem anemic by comparison: like an exercise in giving up on what we should really want, if we could only get it, and *settling* instead for something much less satisfying.

Complicating my tale here is the inconvenient matter that Putnam himself seemed uninterested in any refutation of—or at first even any response at all to—skepticism. Putnam did famously claim to be offering 'an argument we can give that shows we are not brains in a vat' (1981: 8). But curiously, in 'Brains in a Vat,' Putnam at least reads as if he is uninterested in using this argument to combat skepticism *in any way*. Early on, Putnam writes this of the BIV scenario:

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

² See (Williamson 2000b) for critical discussion.

world in a modern way. (How do you know that you aren't in this predicament?) But this predicament is also a useful device for raising issues about the mind/world relationship. (1981: 6)

And in what follows in the chapter, Putnam's own interest seems confined to 'raising issues about the mind/world relationship,' for the 'classical problem of scepticism' is not mentioned again.

Understandably enough, however, others seized on Putnam's argument, or closely related variants of it, as anti-skeptical weapons. And in later comments (Putnam 1994: esp. 284-5), Putnam indicated that his argument was indeed intended to be used against skepticism, but a particular variety of it: 'It was against internal skepticism that my brain-in-a-vat argument was directed' (1994: 284). This 'internal skepticism' is one that aims 'to convince us, on the basis of assumptions we ourselves hold, that all or a large part of our claims about the empirical world cannot amount to knowledge' (1994: 284). What is important to our current concern (determining whether Putnam was trying to refute skepticism in our sense) is what Putnam thought he could avail himself of in battling this 'internal skeptic,' and on this matter, so far from limiting himself to the extremely limited sorts of starting points that avoid begging the question against the skeptic, he seems to embrace an anti-skeptical methodology on which we can use any beliefs we have that are at odds with the skeptic's hypothesis. Eliding over the details about reference that Putnam wields against the BIV hypothesis, here is how Putnam reconstructs his reply:

The internal skeptic I imagined argues that on the basis of our own beliefs about the brain, etc., it follows that we might all be brains in a vat. My reply has the following form. (1) I argue that many of us—perhaps most of us, nowadays—believe that there are causal constraints upon reference ... (2) If we do accept this much about the nature of reference, then the internal skeptic cannot, in fact, show on the basis of premises we accept that we may be brains in a vat. (Putnam 1994: 284-5)

Note that the reason Putnam seems to think that 'the causal constraints upon reference' that he will wield against the skeptic are available for such use seems to be simply that these are things we believe or accept (Putnam seems to intend to use those terms interchangeably), with no mention of their being beliefs of some very narrow and special class that can be legitimately wielded in refutations of skepticism. (It is possible, for all I can really tell, that, though he does not say so,

Putnam is guided here in the premises he chooses to use by aversions to begging the question against the skeptic or to using premises that seem ‘undermined’ by the skeptic’s position, or something like that. If so, then Putnam may be after all seeking something like what we are calling a ‘refutation’ of the skeptic. But if that is going on, it’s not explicit.) Putnam’s thinking here seems to be that since the ‘internal skeptic’ is trying to argue that we may be BIVs from starting points we ourselves accept (‘our own beliefs about the brain, etc.’), it is legitimate in battling this skeptic to appeal to other things we believe in showing that the skeptic’s hypothesis is false. Deriving that the skeptic’s hypothesis is false from other claims that we accept would show that this skeptic does not succeed in her attempt to establish that we may be BIVs in a debate where both sides can take as their starting points the things we ourselves accept.³ The skeptic can appeal to things we accept (primarily, it seems, claims about how our external world beliefs are formed on which it is by a process that could be ‘hot-wired,’ as it were, to form in us perceptual beliefs about the world that are systematically false) to derive the conclusion that we may be BIVs, and this possibility stands unless and until the skeptic’s opponent can show that other beliefs we hold establish that we really aren’t BIVs after all, and that so, preceding from such a starting basis, the conclusion to draw in the end is that, relative to the beliefs we ourselves hold, we really can’t be BIVs.

However, the argument Putnam gives against the BIV hypothesis in (Putnam 1981) does not fit well with the methodological approach he seems to suggest for that very argument at (Putnam 1994). An important feature of the argument itself—and I’m sure a feature in virtue of which it caught and held the attention of so many philosophers—is that it utilizes just the types of premises that might be thought to be legitimate in attempts to *refute* skepticism, in that they might be thought to not beg the question against the prominent form of skepticism many have in mind. As

³ While there are important differences, Putnam’s explanation of what he is up to and why seems to be along lines *fairly* similar to those I later pressed in the closing section of SSP: that since what the interesting skeptic is doing, as I put it, is ‘marshaling deeply felt intuitions of ours in an attempt to give us good reasons for accepting his skepticism,’ it is then of course ‘legitimate to point out that other of our beliefs militate against his position, and ask why we should give credence to just those that favor [the skeptic].’ Indeed, though Putnam does not advocate the whole methodology I follow in SSP, complete with attempting to ‘explain away’ the plausibility of the skeptic’s case, the basic approach he articulates at (Putnam 1994: 284-5) is close enough in content and in spirit to the one I follow that I would be hopeful that he would be quite receptive to the way I have unpacked that basic approach.

we saw in the above quotation, Putnam uses ‘causal constraints upon reference’ to battle the skeptic. But since these constraints are arrived at via thought experiments (in which you make judgments about what a speaker or thinker is referring to by a term in an imagined situation), rather than through empirical investigation of the world, they seem to be candidates for being the kind of *a priori* claims that do not beg the question against the external world skeptic. As we will see in the next section, in what we call the ‘Compatibilist Argument,’ Putnam can be read as also relying on claims about what he is thinking—which are just the sort of ‘internal’ claims about what is transpiring in one’s own mind that form the other main class of premises that are thought to be legitimate as starting points in refutations of the external world skeptic. However, while Putnam’s starting points seem to be carefully chosen so as to allow his argument, if successful, to constitute a refutation of skepticism, the methodology he seems to explicitly advocate in (1994) is not nearly as picky about good starting points, but rather allows one to utilize any beliefs we hold that are at odds with the skeptic’s hypothesis.

The big problem with (what at least seems to be) Putnam’s own reading of his former self is spotted well by Michael Jacovides, whose reading of Putnam’s proposed methodology is similar to mine,⁴ but who then points out that, given such a generous construal of legitimate starting points as that methodology proposes, Putnam’s argument against the skeptic ‘is much subtler than it needs to be’ (Jacovides 2007: 609). For there are many things that we believe much more firmly than the fairly esoteric claims about reference that Putnam uses, but that conflict with the skeptic’s hypothesis that we’re bodiless brains in vats every bit as much as do Putnam’s premises. For example: We believe that we have hands. Indeed (going

⁴ Jacovides writes: ‘If I understand Putnam rightly, the internal sceptic is supposed to be refuted if we can show that something we believed prior to exposure to scepticism is incompatible with the sceptical hypothesis. Since the causal theory of reference is incompatible with it, epistemic progress has been made’ (Jacovides 2007: 609). (Note that Jacovides is not using ‘refute’ as we are here. As we are using the term, Putnam is proposing dealing with this skeptic in a way that does not involve ‘refuting’ her at all.) I found the inclusion of the clause, ‘prior to exposure to scepticism’ in Jacovides’s specification of what Putnam is proposing a bit puzzling, but Jacovides informs me that he did *not* include it in order to rule *out* the use of beliefs in propositions we may not have considered or accepted prior in time to our exposure to skepticism, but rather to rule *in* as acceptable things we used to believe even if our encounter with the skeptic’s arguments may have caused us to lose our beliefs about those matters. At any rate, the important similarity is that on both of our readings of Putnam’s proposed methodology, one does not have to limit one’s starting points in the way needed for one’s argument to constitute a ‘refutation’ of skepticism, as I am using that term.

here beyond points Jacovides makes), given Putnam's apparently very generous construal of legitimate starting points, we may not even need any argument against the BIV hypothesis at all to reject it, for when we hear this hypothesis, most of us believe pretty firmly that it is false. That we are not BIVs would itself then seem to be an available starting point (as well as an ending point) against the skeptic.

Given Putnam's methodological remarks at (1994: 284-5), then, we would expect him to deal with the skeptic very differently, probably not even offering an argument against the BIV hypothesis, but pitting his belief that he is not a BIV itself directly against the skeptic's contention that we may be BIVs or the premises of the skeptic's argument for that contention. Or even if Putnam did offer an argument that he is not a BIV, one would expect it to proceed in a more Moore-like fashion from claims like that he has hands or that he has a body (beyond just a brain), which are intuitively more powerful than are claims about reference from which Putnam actually works, and are also even better candidates for being among 'our own beliefs' or things 'we accept'.⁵

I think it's clear that the interest Putnam's argument has generated does not derive from its being a good one to use when following the kind of methodology Putnam himself seems to be advocating for it at (1994: 284-5). For reasons we just considered, that methodology would lead the wise anti-skeptic who follows it to proceed in a way very different from Putnam's. Rather, the interest in Putnam's argument stems from the fact that it appears to be an attempt to *refute* the skeptic: to show that the skeptic's hypothesis is false from just the sort of argumentative starting points that do not beg the question against her. So, despite Putnam's own explanation of the project he was engaged in, I will here treat Putnam-style arguments in the way that I think they are best viewed: as attempts to refute the skeptic, in the natural sense of 'refute' we are using here. We begin with a quick look at the basic forms such an argument might take.

⁵ A possible reason for using an argument from the likes of *I have hands* or perhaps *I have a body* to the conclusion, *I am not a BIV*, rather than simply taking that conclusion itself instead as a starting point, is that one might think those premises are intuitively more certain than is the proposed conclusion.

4. TWO FORMS OF ARGUMENTS FROM SEMANTIC EXTERNALISM

Putnam-style 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, Putnam suggests, 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 she 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 (1981: 14). Importantly, on any of these suggestions, the BIV would seem to end up thinking something true when it thinks 'There's a tree', or 'Here's a hand', or even 'I'm not a BIV',⁶ 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.⁷

⁶ This aspect of these responses to skepticism is anticipated in (Bouwsma 1949), which argues that a victim of Descartes's evil genius would not be fooled into holding thoroughly false beliefs about the world, but would in fact end up thinking things that were largely true.

⁷ The first form of argument—or what I'm about to call the 'Dilemma' version—seems closer to most of what Putnam himself 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' (Putnam 1981: 8). 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*,

Since I will pursue criticisms of the arguments that don't depend on the fine points of the arguments, we will not need to pursue various details of their formulation, but for the most part can proceed with fairly general statements of their basic form.

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.⁸

We'll call the second type of argument the 'Compatibilist argument.'⁹ 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.

5. OLD OBJECTIONS: VARIETIES OF SEMANTIC EXTERNALISM AND VARIETIES OF SKEPTICAL HYPOTHESES

To set up my own main criticism of these arguments, which will question the anti-skeptical value they would have even if they did in some sense 'work', it will be helpful to first review and then assess a couple of the objections that have already

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 what I'm calling the Compatibilist Argument.

⁸ Readers will have noticed a crucial 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 some of the controversy about the Dilemma argument is over the legitimacy of the needed disquotation. See (Forbes 1995: 212-17) for a good attempt to get from what the argument seems to deliver (*'I am not a BIV' is true*) to what is desired (*I am not a BIV*).

⁹ The argument is so-called because it is available to those who think that 'privileged self-knowledge' of the contents of one's own thoughts is compatible with externalism about the contents of those thoughts.

been leveled against them in the literature, that question whether the argument really does work, 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 to constitute a problem for the Compatibilist form of the argument as well (at least insofar as that Compatibilist argument is supposed to be an effective weapon against skepticism). The objection is that it’s far from clear that a form of semantic externalism that can fuel the anti-skeptical argument is correct.

There *is* strong intuitive 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 himself (1975) and Tyler Burge (1979; 1982; 1986)—in which it’s fairly clear, at least to many 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 anti-skeptical arguments we’re considering. What does this claim tell us about the content of the thoughts of Putnam’s 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 (given that we’re not BIVs), we may well suspect it quite likely indeed that their relevant thoughts differ in content from the analogous thoughts we have, and that the difference in content may be dramatic, 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—or to anticipate a worry of Putnam’s that we will soon note, what *if anything* they 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 anti-skeptical 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 which 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 those involved in 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 very well worth noting here that Putnam himself at times 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 advantages of the Compatibilist form of the argument over the Dilemma form is that the Compatibilist argument does not 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

¹⁰ (Putnam 1981: 15); emphasis added in both quotations.

do BIVs refer to water by 'water'. What's more, if Putnam's occasionally implied worry that BIVs may not succeed in meaning anything at all should prove to be well-founded, then 2 is still on solid ground. Indeed, 2 is true on any account of what if anything the BIVs can mean by 'water' other than that they mean water. And, 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 still seems the kind of premise that would likely be accepted by those who go 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 important second problem for the Putnam-style anti-skeptical strategy:¹¹ that it works at most only against some, but not all, of the ways the BIV hypothesis can be constructed. 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' (Putnam 1981: 6). But what of other scenarios? Putnam's version of the BIV hypothesis seems to be, in David Christensen's nice words, 'cooked up to be vulnerable to the semantical reply' (Christensen 1993: 302). If we instead 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 anti-skeptical 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 seems quite unnecessary, we can go even further, and imagine a scenario of very recent

¹¹ In an earlier paper, I wrote that this problem with the anti-skeptical strategy was 'first noted' (DeRose 1999: 9) in (Brueckner 1986a). 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 counter-attack. (Smith 1984) may be the first place these things were done, but I don't know.

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 many externalists, it will seem that, even now, while briefly 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 anti-skeptical arguments of both varieties.¹² Ted Warfield, in his presentation of his version of the Compatibilist argument, seeks, if somewhat half-heartedly, 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 anti-skeptical arguments, I would like to suggest in closing that we should not be too quick to conclude that the range of applicability of semantic anti-skeptical 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 anti-skeptical arguments depends on how these sorts of issues are resolved and I see no way of resolving them without serious and detailed work on psycho-semantic theories. The common deflationary view of Putnamian anti-skepticism is therefore at least a bit premature. (Warfield 1998: 142-143)

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 anti-skeptical argument from externalism. And he's right that it's far from clear just what range of hypotheses will fall prey to the strategy.

¹² Of course, it is no problem for Putnam himself, if he does not intend to use his argument as an anti-skeptical weapon. If his interest in such an argument only concerns what its availability shows about 'the mind/world relationship,' then he can legitimately be interested in whether a certain version of the BIV hypothesis can be argued against, even if other versions cannot.

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, 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 anti-skeptical strategy and actually suppose that the range of skeptical hypotheses that *escape* being shot down is extremely narrow. *Still the skeptic wins*, for the number of versions of the hypothesis that must work for her to establish her skeptical conclusion is precisely: one.¹³ The status of the war between the skeptic and the Putnam-style anti-skeptic should then be measured by the status of the battle between them over the ways of construing the skeptic’s hypothesis that are most favorable to the skeptic’s case. We should of course evaluate her argument in its strongest form, not in some needlessly weak way.

Putnam seems to have been wise, at least in (Putnam 1981), in not promoting his argument as an anti-skeptical weapon!

6. THE DISADVANTAGES OF HEROISM

As readers may have gathered by now, I find Putnam-style responses to skepticism hopeless—and for well-worn reasons. My purpose in rehearsing the old problems in the previous section, and now moving into new criticisms, is not to somehow make the case against the strategy more decisive: To my thinking my predecessors have already made that case conclusively, far above my poor power to add or detract. But

¹³ 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 piece of knowledge in question. And 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 anti-skeptical result. Still, we want a defense that protects our more detailed knowledge of the external world—like that we have hands—don’t we?

for the purpose of deciding what kind of response to skepticism is most desirable, the issue of what would have been gained if a Putnam-style refutation of BIV-inspired skepticism had been achieved remains important. As I've noted, in attempting to prove that one is not a BIV, the Putnam-style response to skepticism can seem a very aggressive and robust anti-skeptical strategy. Indeed, when a skeptic suggests that you don't know what you think you know because you don't know that you're not a BIV, what stronger response *could* there be than to *prove* to her that you are no such thing? So such a proof/refutation can seem to be the holy grail of anti-skepticism, and attempts to deal with skepticism in other ways can seem disappointing by comparison. But I will now pursue a line of objection 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 anti-skeptical value the Putnam-style arguments would have *even if they did work*. I will argue that in important ways, such a strategy, even if successful, would not provide what we should be looking for in a response to skepticism.

By proving that one is not a BIV, and thereby refuting the skeptic, one seeks, it seems, to *make* it the case that the first premise of AI is false, as applied to oneself. In this respect, this strategy is, as we may say, 'heroic,' in the way Descartes's response to 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.¹⁴ Presumably, though, in either case, the proof provides such knowledge only to 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? Here we encounter what we might think of as the problem of my Mom—whom I use as an example, because, of the people who have not encountered the relevant esoteric philosophy, she, as the person who brought

¹⁴ Thus, after presenting a Putnam-style argument, Colin McGinn writes (one presumes with a good deal of irony): 'In effect, I can achieve the anti-sceptical result Descartes needed God to vouchsafe by exploiting considerations about what determines content. Ah, the wonders of analytical philosophy' (McGinn 1989: 113). (Jacovides 2007) argues that in an important way, Descartes's anti-skeptical arguments are better than Putnam's. I should note, though, that I don't think that Descartes was seeking to gain or defend *knowledge* by his proof, but some higher state: see fn. 20 of Chapter 4, and, for a fuller account, (DeRose 1992b).

me into the world, would perhaps be the person it would be worst for me to ‘leave behind’, in a state of not knowing even that she has hands!¹⁵

So: What of my Mom? Does *she* fail to know that she has hands? Our skeptical argument could be applied to her, after all.¹⁶ Non-heroic strategies, which attempt not to show how to gain knowledge in the face of the skeptical argument, but rather to show how—or at least to make the case that—the skeptical argument *never worked in the first place*, by protecting the knowledge of non-philosophers, 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 Putnam-style 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:

¹⁵ My Mom originally entered the discussion when I was working on the talk that became (DeRose 2000a), which forms the basis for much of this chapter. I used her as my example for the reason just given in the text above, but, since I sent her the paper in which she was used as an example, she has ceased to be a good example, since she has now encountered the relevant arguments. (I don’t think she accepts them—but she might just be trying not to hurt my feelings!) But we can still use her past self as an example.

¹⁶ AI can be applied most smoothly to those who haven’t considered the anti-skeptical ‘proofs’ when we suppose that they have encountered the skeptical hypothesis itself, especially if they have encountered it as part of a skeptical challenge. In that case, we can get a version of AI that’s about as powerful as the original just by replacing ‘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.

Things get trickier when we consider those who have not even encountered the skeptical hypothesis. So, suppose that is true of my Mom. We may then worry that she 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. See Appendix D for how to apply AI to those who have not even considered the skeptical hypothesis. For our discussion here, we will just suppose that my Mom has considered the skeptical hypothesis, and perhaps the skeptical argument in which it is featured, but has not encountered any proposed fancy philosophical arguments designed to show that it does not obtain.

AI:

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 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 Putnam-style argument to the conclusion that I'm not a BIV help in this predicament?

AI presents a very limited range of opportunities for a proof of not-H to help.¹⁷ That there is such an 'anti-skeptical' argument, even should it be sound, and

¹⁷ Well, AI *itself* doesn't make much room for a proof of not-H to help, other than by making AI's first premise go false, if need be. But we could certainly work out, say, a good sense of 'possible'—we could then label it 'possible_p'—on which H is not 'possible' if there is a Putnam-style proof (a proof proceeding from premises of a favored sort) that it is false. We won't bother here to work out the exact meaning of 'possible_p', or the favored sort of premises in question, because the important point to follow should apply however exactly that is all done. A Putnam-style proof of not-H then would rebut a premise of a modified AI, which has a third premise added to it, like this: :

AI+:

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

But that added third premise is completely unnecessary. We could make it necessary by weakening the second premise, yielding:

AI+2:

1. I don't know that not-H
 - 2w. If I don't know that not-H, and if H is possible_p, then I don't know that O
 3. H is possible_p
- So,C. I don't know that O.

Though he does not really spell out how the argument is supposed to go, Nozick seems to have in mind an argument something like this, complete with something like a premise of possibility ('logical possibility' seems the variety he has in mind), especially at (Nozick 1981: 167-9), where he's setting up his chapter on 'Knowledge and Skepticism.' And Nozick raises Putnam's response at this point (see the note that runs at the bottom of pp. 168-9). But by this point we seem not to be searching for a good skeptical argument so much as an argument that is (to echo Christensen's words) cooked up especially for the purpose of being vulnerable to

sound in a way that allows it to constitute a refuting proof of its conclusion, doesn't seem to have any tendency to show that AI is invalid, or that its second premise is false, or that its first premise wasn't true of me before I knew of the anti-skeptical argument. The only readily apparent way the anti-skeptical argument could help with AI itself is by *giving* me knowledge, should I not already have it, that not-H, and thereby making the first premise come to be false, as applied to me, if it isn't already false. And those who discuss the semantic anti-skeptical 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

a Putnam-style attack. Original, unweakened 2 seemed quite plausible already. I suppose one might think that 2's plausibility depends on our presuming that H is possible_p, for some good working out of 'possible_p', so that AI+2 really does make explicit an important ingredient in the power of AI-like arguments. Such thinking would likely be connected with thoughts to the effect that the 2-like premise in such an argument should be the instance of a good formulation of the closure principle, together with some ideas of what such a principle would look like. For my part, I am extremely open to the need to tweak the closure principle in all sorts of ways, but don't think that modifications that would result in an AI-like argument that would make a premise like 3 needed are among those that would really help. For my thoughts on the ways closure should be tweaked, and the advanced form of the skeptical argument that then arises, see sections 20-22 of Chapter 5 and Appendix D.

¹⁸ Anthony Brueckner's treatment of the strategy in his (1992) 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' (Brueckner 1992: 200)—our premise 1—and it's this premise which he says is targeted by the semantic anti-skeptical strategy, which is described simply as 'blocking the skeptical argument' (1992: 202), apparently at the premise we've identified. Brueckner makes some of the preliminary moves of the anti-skeptical argument, and then asks the crucial question: 'How does any of this help with the refutation of Cartesian skepticism?' (1992: 204), but then just dives into the details of the argument, without explaining how such an anti-skeptical 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' (1992: 205), he writes, 'So if we accept the externalist semantic assumptions underlying the argument, it seems that we do get the desired anti-skeptical conclusion that I am not a BIV' (1992: 205). 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.

So I'll here assume that the semantic anti-skeptical strategy is supposed to help us with AI by heroically giving us knowledge of not-H (if we don't already have it), and thereby making AI's first premise false as it is applied to us (if it isn't false already). And if that is how the anti-skeptical argument helps, then, of course, as I've charged, it doesn't help at all in protecting the knowledge of those who have never grasped the argument.

And the problem doesn't stop there. 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. And even among those who promote this anti-skeptical 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. (And since this broad kind of destructive interaction can happen with the problems with other heroic attempts to refute skepticism, it's worth briefly pursuing the issue.) 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 the 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, in order to protect someone's knowledge that they have hands from the skeptical argument, it's not enough for the anti-skeptical argument to work against just some versions of the BIV hypothesis, but it rather has to work against even, say, the recent envatment scenarios, then, if the argument only works for those who gain knowledge by means of grasping it, there is a very substantial chance, in my estimation, that *even if the anti-skeptical argument against recent envatment is sound* (which it almost surely isn't), there is not a single human being 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 exceptionally safe 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 anti-skeptical argument is extremely narrow.

Of course, for all I've argued, if the Putnam-style arguments are in fact sound proofs, and are sound proofs as applied to all the hypotheses they'd have to work against, then there is some (but perhaps 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 and quite important worry concerning the heroic nature of this response is worth mentioning: 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 ordinary call of knowledge.¹⁹ 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. And 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

¹⁹ 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 to satisfy the standards for knowledge set in most ordinary contexts.

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.²⁰

7. THE CHALLENGES FACING NON-HEROIC, MOOREAN ALTERNATIVE RESPONSES

So it seems we really do need a non-heroic 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—i.e., one according to which it was never the case that both of the skeptical argument's premises were true of us. Or in the case of contextualist responses: one on which it was at least never the case that both of the skeptical argument's premises were *simply* true of us, nor that they were ever both true of us when evaluated at ordinary standards. Recall, however, that we are in this chapter considering only (substantively) Moorean responses to AI, and so have set aside responses that deny the second premise of the skeptical argument. (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!)

²⁰ I for one think philosophical arguments quite generally (even if perhaps not exceptionlessly) fail to yield knowledge of their conclusions: see Appendix B.

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 many of 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 doesn't clear the bar for knowledge by very much. 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. And 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, pre-philosophical 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.

Non-heroes who accept the skeptic's second premise seem to have three options—the good, the bad, and the ugly. We'll canvass them quickly in reverse order. First, they can accept that our knowledge of such propositions as that we have hands is at best 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 noted in the closing paragraph of the preceding section, 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. 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 really depends on the latter in the way envisioned.

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. And recall, this is to be a non-heroic 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 pre-philosophical, 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. Unsurprisingly, then, a satisfying non-heroic response (one according to which the skeptical argument was never sound—or at least was never *simply* sound—in the first place) will have to be what we may call a *diagnostic* response: one which *explains how we were misled by some component of the argument*. And in this case, it should not only explain how the argument's first premise is wrong, but also how it is so far wrong: how it might seem true, while we in fact not only know that we are not BIVs, but know this in a very solid manner.

8. COMPARING THE TWO MOOREAN RESPONSES TO SKEPTICISM

And of course, my response to AI in SSP is such a diagnostic response. Realizing that responses to skepticism that don't even seek to refute the skeptic can seem feeble, my purpose here has been to show how a *mere* refutation of skepticism (as we might now put things) would leave a lot to be desired. In particular, it would not adequately address the standing of those folks who do not possess the refutation. This combined with how such a diagnostic response can help us in the philosophical pursuit of responding rationally to the skeptical argument in question, in the way explained in Chapter 2, constitutes my defense of the type of response I give.

In brief summary, then, the two responses to AI we have been considering in this chapter respond very differently to the apparent shakiness of whatever belief we might have to the effect that we're not BIVs, the would-be externalist refuter of

skepticism 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, and we have this knowledge before engaging in any philosophy. In chapter 7, I will explain how.