

A MOOREAN RESPONSE TO BRAIN-IN-A-VAT SKEPTICISM

Australasian Journal of Philosophy 80 (2002): 148-163

Tim Black

Certain deceptively simple skeptical arguments threaten our knowledge of the external world. The following is an example of such an argument.

The Argument from Ignorance (AI)

1. I don't know that I'm not a brain-in-a-vat (henceforth a BIV).
2. If I know that I have hands, then I know that I'm not a BIV.
3. Therefore, I don't know that I have hands.¹

We can respond to AI in at least three ways. First, we can accept it outright. Call this *the skeptical response*. Second, we can respond as G. E. Moore might. Moore would probably accept AI's second premise, but then claim that he knows that he has hands. From this it would follow that he knows that he is not a BIV. This is *the Moorean response*. Third, we might be able to reconcile these two responses. One way to do this would be to show that we should accept the Moorean response in some contexts but the skeptical response in others. Keith DeRose's recent *contextualist response* fits into this third way of responding to AI.²

In what follows, I argue that we can provide a Moorean invariantist response to AI, a response according to which the standards for knowledge are *always* comparatively *low*.³ I also

¹ See DeRose [8], p. 183.

² DeRose introduces his contextualism in [5] and presents it in further detail in [8]. He also defends his view against important objections in [6] and [7].

³ When I say that I will provide a Moorean response to skepticism, I do not mean to suggest that my response is *Moore's* response. There are differences as well as similarities between the two. The main – and perhaps the only – similarity between my response and Moore's is this: both responses claim that the standards for

argue that this Moorean response has all of the advantages and none of the disadvantages of DeRose's contextualist response. The Moorean response that I provide explains why we make the epistemic judgements that we make, and it provides these explanations while employing the same theoretical framework that DeRose employs, namely, Robert Nozick's subjunctive conditionals analysis of knowledge.⁴ Furthermore, my Moorean response avoids two problems that confront DeRose's contextualism. I contend, therefore, that we should prefer my Moorean response to DeRose's contextualist one.

Before I proceed, I should say that I respond in the present paper only to AI, which is based on the BIV skeptical hypothesis. It is true, of course, that we can build skeptical arguments from different skeptical hypotheses. We can build skeptical arguments from the hypothesis that I might now be dreaming, or from the hypothesis that I am being deceived by an omnipotent and malevolent god. I concentrate on AI because it is one of the most troubling skeptical arguments around, if not *the* most troubling. We will thus have accomplished a great deal if we can provide an adequate Moorean response to AI. Moreover, even if we discover arguments that are more troubling than AI, I believe that my response to the BIV skeptical argument is an instructive first step towards a decisive Moorean response to skepticism. I will have more to say about this in Section IV.

I. Trouble with contextualism

According to DeRose, we tend to judge in most contexts that I know that I have hands. But he also believes that we tend to judge in other contexts that I *don't* know that I have hands. His primary argument for the contextualist response to skepticism rests on the claim that contextualism

knowledge are low, low enough so that I can know both that I have hands and that certain skeptical hypotheses do not obtain (e.g., that I'm not a BIV). But there are also significant differences between the two responses. For example, while the notions of certainty and doubt play essential roles in Moore's response to skepticism, neither notion plays a role in my response.

best explains why we make these judgements. Contextualism says that the standards for knowledge shift from context to context. In some contexts, the standards for knowledge are unusually high, and the skeptic can truthfully say in those contexts that I *don't* know that I have hands. Contextualists like DeRose suggest that we should therefore accept the skeptical response in those contexts. However, the standards for knowledge are comparatively low in other contexts, and thus we can correctly say in those contexts that I *do* know that I have hands. We may therefore accept the Moorean response in those contexts.⁵ DeRose claims that his contextualist response to skeptical arguments best explains why we make the epistemic judgements that we do, and thus that we should prefer his contextualist response over both the skeptical response and the Moorean one.

I want briefly to address two difficulties that confront DeRose's contextualist response to skepticism. These difficulties do not confront the Moorean response that I provide later. First, at the heart of DeRose's contextualism – indeed, at the heart of many other extant contextualisms⁶ – is the claim that the standards for knowledge shift from context to context. In particular, contextualists claim that the skeptic can *raise* those standards. Of course, this gives rise to the question *how* might the skeptic raise the standards. Contextualists have provided a variety of answers to this question. DeRose suggests that the skeptic can raise the standards by asserting that I don't know that I'm not

⁴ See Nozick [23].

⁵ DeRose and other contemporary contextualists are primarily concerned with the truth conditions of knowledge-ascribing and knowledge-denying sentences, sentences such as 'Keith knows that he has hands' and 'Stewart doesn't know that he's not a BIV' (see, for example, DeRose [6]). Nevertheless, these contextualists often 'semantically descend' from talk of the truth conditions of such sentences to talk of the standards for knowledge (see DeRose [7]). DeRose seems to believe that such a semantic descent is permissible as long as we remember that the primary concern is with the truth conditions of certain sentences. So, like DeRose and others, I often talk of the standards for knowledge rather than of the truth conditions of knowledge-ascribing and knowledge-denying sentences.

Also, for DeRose, as for several other contemporary contextualists, *contexts* seem to be *conversations*. So whatever the criteria are for the individuation of conversations, those same criteria, or at least some very similar ones, will be the criteria for the individuation of contexts.

⁶ These other contextualisms include those advocated by Cohen [2], [3], Heller [17], Lewis [21], and Rieber [25].

a BIV.⁷ The skeptic's assertion invokes a conversational rule that DeRose calls the Rule of Sensitivity:

When someone asserts that S knows (or does not know) that P, the standards for knowledge tend to be raised, if need be, to a level such that S's belief that P must be sensitive if it is to count as knowledge.⁸

But some contextualists claim that the invocation of certain conversational rules isn't, or isn't always, enough to raise the standards for knowledge. For it seems that the skeptic's conversational partners often successfully resist the raising of the standards. Thus, some contextualists suggest that the skeptic's conversational partners must *cooperate* with the skeptic, allowing her to raise the standards for knowledge, if she is to raise those standards.⁹ It seems, however, that there are contexts in which it is the case both that the skeptic's conversational partners are perfectly cooperative and that she nevertheless fails to raise the standards. For example, the skeptic's cooperative conversational partners might unwittingly resist the raising of the standards by failing to realize that she is trying to raise them. In light of difficulties like these, Stewart Cohen has recently suggested that 'the standards are determined by a complicated pattern of interaction among the intentions, expectations, and presuppositions of the members of the conversational context'.¹⁰

It seems, then, that there is no consensus among contextualists with respect to an explanation of how the skeptic might raise the standards for knowledge. Thus, they fail to agree on

⁷ Although DeRose does not subscribe to the view that the skeptic's assertion raises the standards for knowledge *without fail*, some contextualists do seem to subscribe to this view. See, for example, Cohen [2], p. 108, and Rieber [25], pp. 196-197.

⁸ DeRose [8], p. 206. My belief that P is *sensitive* if I wouldn't believe that P if P were false. Thus, to say that my belief that P must be sensitive in order to count as knowledge is to say that I know that P only if I wouldn't believe that P if P were false. My belief that P is *insensitive* if I would believe that P even if P were false.

⁹ Contextualists who suggest that cooperation is important (or even necessary) for raising the standards include Lewis [19], Heller [17], and Cohen [3], especially p. 61 and p. 88, n. 63. DeRose too suggests that cooperation might be important (see [8], p. 186), but he does not emphasize it to the extent that Lewis, Heller, and Cohen do.

¹⁰ Cohen [3], p. 88, n. 63.

an explanation of a claim that lies at the very heart of their project. Moreover, even if contextualists were to agree to opt for a line like Cohen's, they still would not have provided a precise explanation of how shifts in the standards for knowledge occur, for Cohen's explanation does not specify the intentions, expectations, and presuppositions involved, or how these things must interact in order to bring about a shift in the standards. The lack of such a specification might lead us to wonder whether the 'complicated pattern of interaction' can ever really occur and thus whether the standards for knowledge can ever really shift. And since the lack of such a specification casts doubt on these essential claims, it leaves contextualism vulnerable to alternative responses to skepticism, especially if those alternatives have all of contextualism's advantages but do not demand, or even compel us to try to provide, an explanation of how the standards for knowledge can shift from context to context. My Moorean response, or so I shall argue, is just such an alternative. We should therefore prefer it over contextualism.¹¹

There is a second difficulty with DeRose's contextualist response to skepticism – it forces him to admit that there are contexts in which we should accept the skeptical response to BIV arguments. Yet even this small concession to the skeptic – if it is a *small* concession – is somewhat disagreeable, if not altogether unpalatable. Furthermore, if we find these concessions to the skeptic hard to swallow, perhaps we should favor a Moorean response. For a Moorean response, unlike either the skeptical response or DeRose's contextualist response, does not force us to make any concessions to the skeptic.

¹¹ DeRose has suggested, in correspondence, that the task of providing an accurate specification of the conditions under which the skeptic might raise the standards for knowledge is a difficult one. He also suggests, however, that we should not impeach the contextualist response to skepticism simply because contextualists have not yet provided such a specification. (For similar suggestions, see Cohen [1], p. 24, n. 19, and Cohen [4], p. 98.) Yet, once again, if there is a response to brain-in-a-vat skepticism that has the advantages and lacks the disadvantages of contextualism and that does not call for a complicated and forthcoming (at best) explanation, then we should opt for that response.

So the Moorean response that I propose faces neither of the two difficulties that confront contextualism. Since it is a Moorean *invariantist* response – that is, since it is a response according to which the standards for knowledge are *always* comparatively low – it is a response according to which the skeptic cannot raise the standards for knowledge. So we need not explain how the skeptic might raise those standards. Furthermore, since the response that I propose is a *Moorean invariantist* response, it does not force us to admit that there are contexts in which we should accept skepticism. Thus, the response that I provide has neither of contextualism’s disadvantages. It does, however, have all of its advantages. In Section III, I argue that the Moorean response provides an explanation of our epistemic judgements that is just as good as the explanation provided by contextualism. This is not to suggest that there is no work that needs to be done on behalf of the Moorean response. For it commits us to the idea that I can know across contexts that I’m not a BIV, and this claim strikes some as implausible. So the Moorean response raises questions concerning how I can know that I’m not a BIV and questions concerning why we tend to judge in some contexts that I *don’t* know that I’m not a BIV. Fortunately, we can answer these questions. Thus, like the contextualist response, the Moorean response provides an adequate explanation of our epistemic judgements. When we put this together with the fact that it confronts neither of the difficulties that confront contextualism, we have good reason to prefer the Moorean response over DeRose’s contextualist response.

II. Nozick’s account of knowledge

DeRose’s contextualist response to skepticism is based on Robert Nozick’s subjunctive conditionals account of knowledge. Yet if we can use Nozick’s account in providing a Moorean response to skepticism – and we will see that a certain underappreciated aspect of the subjunctive conditionals account allows us to do so – then our Moorean response will be based on a theoretical foundation that DeRose should find uncontroversial. So he will not be able to say that our response

to skepticism and the explanations it provides are inadequate because they are based on an objectionable theoretical foundation.¹²

Nozick says that S knows that p if and only if the following four conditions are met: (1) S believes that p; (2) p is true; (3) S wouldn't believe that p if p weren't true; and (4) S would believe that p if p were true. There is some controversy concerning how we should evaluate the subjunctive conditionals in conditions (3) and (4). Some suggest that only the *closest* possible worlds – only those possible worlds that are most similar to the actual world – are relevant to their evaluation,¹³ while others suggest that there are more distant possible worlds that are close *enough* to the actual world to be relevant.¹⁴ I believe that there are persuasive objections to the former suggestion.¹⁵

¹² Those who do object to Nozick's account might object to my response on the grounds that it is based on that account. Fortunately, though, my response does not depend on Nozick's account. As we shall see, I emphasize certain core features of that account and, in particular, a certain underappreciated feature that concerns our belief-forming methods. Furthermore, the notion of a belief-forming method (or some similar notion) plays a prominent role in any number of other theories, and my response might have been based on virtually any of those theories. (Such theories include those proposed by Goldman [11], [12], [13], [14], Greco [15], Plantinga [24], and Sosa [26].) Moreover, even though none of these theories is without its critics, many agree that a notion similar to that of a belief-forming method must play a starring role in any adequate theory of knowledge. (For an argument for this claim, see Goldman [12], pp. 42-43.) So, by emphasizing the underappreciated feature of Nozick's account, I ensure not only that my response doesn't depend on that account, but also that it hinges on a feature of his account that seems a necessary ingredient in any adequate theory of knowledge. For these reasons, my response to skepticism should be attractive from many theoretical points of view.

¹³ See Goldman [12], pp. 45-46.

¹⁴ See Heller [18], especially pp. 115-116 and p. 128, n. 3.

¹⁵ Goldman ([12], p. 45) has provided the following consideration against the former suggestion. Suppose that there are many thermometers in Sue's medicine cabinet and that all but one are defective. She cannot tell which thermometer is which, but she just happens to pick the reliable thermometer when she suspects that her son has a fever. She uses that thermometer to take her son's temperature, and it correctly reads 98.6 degrees. Sue believes on the basis of this thermometer reading that her son's temperature is normal. On the former suggestion, Nozick's conditions on knowledge are met in this case. Condition (3) is met because Sue wouldn't believe that her son's temperature is normal if it weren't. For the *closest* worlds in which her son's temperature is abnormal are worlds in which she uses the reliable thermometer and are thus worlds in which she believes that his temperature is abnormal. Condition (4) is met because Sue would believe that her son's temperature is normal if it were. For, again, the *closest* worlds in which her son's temperature is normal are worlds in which she uses the reliable thermometer and are thus worlds in which she believes that his temperature is normal. But Sue does not know in this case that her son's temperature is normal, for she cannot tell which thermometer is which and it is just luck that she picked the reliable one. The former suggestion is inadequate because it does not account for certain relevant alternative situations, namely, those in which Sue selects one of the many defective thermometers. These relevant alternative situations are not among the *closest* situations (since the closest situations are ones in which Sue selects the reliable thermometer, the thermometer that she actually selects), but they are among the situations that are close *enough* to affect whether or not Sue knows.

Thus, I opt for the latter suggestion, according to which both the *closest* and the close *enough* worlds are relevant to the evaluation of the subjunctive conditionals in (3) and (4).

In fact, besides allowing us to avoid certain objections, there are other advantages of opting for the latter suggestion. The theoretical apparatus behind that suggestion sustains not only DeRose's contextualism but also the contextualisms of those who are motivated primarily by the relevant alternatives theory of knowledge.¹⁶ Thus, when the latter suggestion is the backdrop for our Moorean response to skepticism, our response can serve as an alternative to a wider range of contextualist responses, a range that includes both DeRose's contextualism and contextualisms that are motivated by the relevant alternatives theory.¹⁷ From now on, then, I will understand Nozick's analysis of knowledge in terms of the latter suggestion.

After initially stating his conditions on knowledge, Nozick restates them in order 'to take explicit account of the ways and methods of arriving at belief'.¹⁸

S knows that p if and only if

(1) p is true;

(2) S believes, via method or way of coming to believe M, that p;

¹⁶ The relevant alternatives theory was first proposed by Dretske [9], [10] and by Goldman [11].

¹⁷ The contextualisms of Cohen [2], Lewis [21], and Heller [17] are each based on the relevant alternatives theory of knowledge. (For a nice account of the similarities between the subjunctive conditionals account of knowledge and the relevant alternatives theory, see Heller [16].) Even though my response to skepticism is expressly designed to be an alternative to contextualisms that are based on externalist theories such as the subjunctive conditionals account and the relevant alternatives theory, it also has much to say against contextualisms that are based on internalist theories (see Cohen [3]). For example, since my Moorean response commits us to no disagreeable skeptical conclusions, we have reason to prefer it over *any* contextualism that does so, including those that are based on internalist theories. Also, the fact that an externalist theory helps us adequately to respond to skepticism might give us some general reason to prefer such a theory over internalist ones, and this might in turn give us reason to eschew internalist contextualisms.

¹⁸ Nozick [23], p. 179. I suspect that the difficulties with *method* individuation run parallel to the difficulties with *process* individuation (see Goldman [14]). Later, I will have a bit to say about method individuation, but for the most part I hope to avoid this difficulty. I rely on the somewhat standard classification of methods – perception (including seeing, hearing, tasting, touching, and smelling), memory, testimony, intuition, and other less highly esteemed methods such as clairvoyance and fortune telling. By emphasizing the importance of methods, we do not lose sight of contextualism. See DeRose [8], pp. 195-196, 204, 206.

(3) If *p* weren't true and *S* were to use *M* to arrive at a belief whether (or not) *p*, then *S* wouldn't believe, via *M*, that *p*; and

(4) If *p* were true and *S* were to use *M* to arrive at a belief whether (or not) *p*, then *S* would believe, via *M*, that *p*.¹⁹

It seems to me that Nozick's revised conditions and their implications have been unduly neglected. I believe, however, that we can utilize them in an important way in providing an adequate response to skepticism.

Nozick provides the following example as support for his revised account: 'A grandmother sees her grandson is well when he comes to visit; but if he were sick or dead, others would tell her he was well to spare her upset'.²⁰ When her grandson is well, the grandmother believes on the basis of seeing him that he is well. But if he were not well, she would use another method – testimony – in coming to believe whether or not her grandson was well. And in that case the belief she would hold would be false. Yet, as Nozick says, the fact that she would use another method of arriving at belief if her grandson were not well 'does not mean that she doesn't know he is well (or at least ambulatory) when she sees him'.²¹ This suggests that the only worlds that are relevant to *S*'s knowing that *p* (where *p* is a proposition about the external world that *S* ordinarily takes herself to

¹⁹ Nozick [23], p. 179.

²⁰ Nozick [23], p. 179.

²¹ Nozick [23], p. 179. Nozick also provides a second example in support of his revised account of knowledge. Suppose that Silas believes in the actual world that it rained in London yesterday. Silas was in London yesterday, he saw then that it was raining there, and he now remembers that it did. In this case, we are inclined to think that Silas knows that it rained in London yesterday. But suppose that there is another world *W* in which Silas has been stuck in New York for three days. In *W*, Silas's friend mistakenly reports to him that it didn't rain in London yesterday. So Silas believes in *W* that it didn't rain in London yesterday. Nozick says that 'this does not show [that Silas] actually doesn't know that [it rained in London yesterday], for actually he has not used this alternative method [namely, testimony] in arriving at his belief' ([23], p. 179).

know) are worlds in which, in arriving at the belief that *p*, *S* uses the same belief-forming methods that she actually uses.²²

I now want to emphasize an important implication of Nozick's revised analysis of knowledge, an implication that, so far as I know, has not yet been fully appreciated by those contextualists who employ Nozick's subjunctive conditionals analysis. As we have just seen, the only worlds that are relevant to whether or not I know that *p* are those in which my belief is produced by the method that actually produces it. This feature of Nozick's analysis has interesting and heretofore unnoticed consequences. Among these is the fact that BIV worlds – possible worlds in which I am a BIV – are *not relevant* to whether or not I know that I'm not a BIV. For, as I later argue, BIV worlds are worlds in which my belief is produced by a method other than the one that actually produces it. Moreover, since BIV worlds are not relevant to whether or not I know things about the external world, we may provide a Moorean response to brain-in-a-vat skepticism, a response according to which I can know both that I have hands and that I'm not a BIV.

III. A Moorean response to AI

Nozick's analysis of knowledge gives us room to provide a Moorean response to brain-in-a-vat skepticism. To show that it does, I must provide three explanations. First, I must explain how AI's conclusion can be false. That is, I must explain how I can know that I have hands. Second, in providing a Moorean response to skepticism, I will claim that AI's first premise can be false. So it is

²² Two complications arise here. First, complications arise when a particular belief is (actually) produced by more than one method. I hope for the most part to ignore such complications, but for a nice treatment of them, see Nozick [23], pp. 179-185. Second, it might be that there are cases in which one method – say, memory – (actually) produces *S*'s belief that *p* but another method – say, testimony – (actually) produces a belief that is incompatible with her belief that *p*. Cases such as these do not adversely affect Nozick's analysis of knowledge. In such cases, *S*'s belief that *p* will fail to satisfy either condition (3) or condition (4). Likewise for the incompatible belief.

perhaps most important that I explain both how I can know that I'm not a BIV and why it is sometimes plausible to say that I *don't* know that I'm not a BIV. Finally, I must explain why AI's second premise is true.

First, how can I know that I have hands? I can know that I have hands because my belief that I have hands sometimes meets Nozick's four conditions on knowledge. In the actual world, as well as in other nearby possible worlds, it is true both that I have hands and that I believe via perception that I have hands. But there are some pretty close worlds in which I don't have hands. These include worlds in which I lost my hands about twelve years ago in an unfortunate forklift accident. Yet in the no-hands worlds that are close enough to the actual world, I don't believe that I have hands if I arrive via perception at a belief as to whether I have hands. So my belief sometimes meets conditions (3) and (4) of Nozick's account and I can thus know that I have hands.

Second, I can know that I'm not a BIV because my belief that I'm not a BIV sometimes meets the conditions on knowledge. First, it is true that I'm not a BIV. Next, I believe that I'm not a BIV either via perception – perhaps I believe that I'm not a BIV because I perceive that I have hands – or via an inference from perceptual beliefs – perhaps I believe that I'm not a BIV on the basis of an inference from my perceptual belief that I have hands.

Furthermore, my belief that I'm not a BIV can meet conditions (3) and (4). To see that this is so, we must first circumscribe the sphere of worlds throughout which either perception or an inference from perceptual knowledge produces my belief as to whether I'm a BIV. Now, there are no BIV worlds in this sphere, for BIV worlds are worlds in which neither perception nor an inference from perceptual knowledge produces my belief. In BIV worlds, my belief that I'm not a BIV is based (either directly or inferentially) on my phenomenal experiences, which are produced by scientists who are using a sophisticated artificial process in order to stimulate my brain electrochemically. And this method is not the method of perception.

Of course, this claim meets with resistance. Nozick says,

Usually, a [belief-forming] method will have a final upshot in experience on which the belief is based, such as a visual experience, and then (a) no method without this upshot is the same method, and (b) any method experientially the same, the same “from the inside”, will count as the same method. Basing our beliefs on experiences, you and I and the person floating in the tank are using ... the same method.²³

But Nozick’s conclusion – the conclusion that I and the BIV use the same belief-forming method – is a bit hasty. Consider two belief-forming methods, A and B. Let A be visual perception, however we are properly to characterize that method, and let B be some belief-forming method. Given this, Nozick suggests the following:

B is visual perception if and only if it is experientially the same, the same ‘from the inside’, as A. That is, B is visual perception if and only if its final experiential upshots are the same as the final experiential upshots of A.²⁴

But there is a problem with this analysis. It counts as visual perception methods that seem to be something other than visual perception.

Suppose that Ray has been blind from birth. One day, however, he has a phenomenal experience as of a purplish light, and he mistakenly believes on the basis of this experience that he is gaining his sight. Of course, Ray’s purplish experience is the final upshot of some method – call it Ray’s belief-forming method – and we may suppose that his experience, which is the only experience of its kind that he has had, is the same ‘from the inside’ as a visual experience. (Perhaps

²³ Nozick [23], pp. 184-185.

²⁴ This principle results from combining Nozick’s (a) and (b). (a) says that no belief-forming methods without visual experiences as final experiential upshots are visual perception. That is, (a) says that X is visual perception *only if* it has visual experiences as final experiential upshots. (b) says that all methods with visual experiences as final experiential upshots are visual perception. That is, (b) says that X is visual perception *if* it has visual experiences as final experiential upshots.

it is the same ‘from the inside’ as my visual experience of a purple light flashing in a dark room.) Now, according to Nozick’s analysis, Ray’s belief-forming method – the method that has as its final upshot Ray’s purplish experience – counts as visual perception. But this seems to be a mistake; ordinarily, we would not say that Ray’s belief-forming method counts as visual perception. Just as we might say that the final upshots of last night’s dream are the same as the final upshots of visual perception, we might say that the final upshot of Ray’s belief-forming method is the same as the final upshots of visual perception. But just as in the case of dreams, we conclude from this only that Ray has had an experience of a certain sort. We do not conclude that Ray’s belief-forming method – or last night’s dreaming – counts as visual perception. This suggests that Nozick’s account of visual perception is unsatisfactory. If a method is to count as visual perception, it must satisfy not only Nozick’s ‘inside’ conditions – being experientially the same, the same ‘from the inside’, as visual perception – but also some additional conditions. And these additional conditions will be ‘outside’ conditions.

This allows us to say, then, as perhaps we should, that the method of visual perception includes

the lenses of the eyes focussing ... light on the retinas, where a pattern of retinal cell stimulation occurs that sends electro-chemical impulses along the optic nerve to the visual cortex, where a pattern of brain cell stimulation occurs with the upshot that the subject has a visual experience.²⁵

²⁵ McLaughlin [22], p. 200. I do not mean to suggest that our notion of visual perception includes *only* these things. Our notion of visual perception seems to include certain core physiological and neurophysiological elements while still allowing for variation. Thus, our notion of visual perception can accommodate the fact that human beings with prosthetic rather than natural eyes see. See Lewis [20].

Now, it seems that we should characterize visual perception in this way – at least in part ‘from the outside’ – if we want to distinguish visual perception from other methods that we do not count as visual perception, for example, Ray’s belief-forming method. Typically, processes like the one described above must occupy a central position in the characterization of a belief-forming method if we are to count that method as visual perception. And since no such process helps to characterize Ray’s belief-forming method, we do not count that method as visual perception. Furthermore, we should also characterize other perceptual belief-forming methods ‘from the outside’ if we want to distinguish them from methods that do not count as perceptual belief-forming methods. This means that things like (properly-functioning) retinas, tympanic membranes, olfactory nerves, and taste buds will help to individuate perceptual belief-forming methods.²⁶ Of course, we can characterize the method of visual perception in this way – at least in part ‘from the outside’ – while still maintaining that visual experiences are the final experiential upshots of visual perception.

Now, in the actual world, my belief that I’m not a BIV is either directly based on the upshots of perceptual processes or inferred from beliefs that are themselves directly based on such upshots. In either case, perceptual processes play a substantial role in producing my belief that I’m not a BIV. But, *ex hypothesi*, BIVs lack retinas, tympanic membranes, olfactory nerves, and taste buds. Thus, since these things help to make our perceptual processes what they are, the BIV’s belief that it isn’t a BIV is produced by methods other than perceptual ones. It follows that no BIV worlds are

²⁶ Now that we have characterized visual perceptual belief-forming methods as we have, it might seem that Nozick’s conditions on knowledge are trivially satisfied, for the use of such methods entails that I’m not a BIV. But although my using the method of visual perception *does* entail that I’m not a BIV, this does not constitute the trivial satisfaction of Nozick’s conditions, for my using that method does not entail that I *believe* via that method that I’m not a BIV, and what is in question in those conditions is what I *believe*.

Also, I will from now on leave out the ‘properly-functioning’ qualification. Nevertheless, unless otherwise stated, I intend for the qualification to apply.

worlds in which my belief that I'm not a BIV is produced by the same methods that produce it in the actual world. Thus, no BIV worlds are relevant to whether or not I know that I'm not a BIV.

We can now see that I can know that I'm not a BIV, for we can now see that the counterfactuals in conditions (3) and (4) are both true for my belief that I'm not a BIV. The counterfactual in condition (3) is true because there are no worlds in which I'm a BIV and in which perception substantially produces my belief as to whether I'm a BIV. The counterfactual in condition (4) is true as well. We have said that if I am to know that I'm not a BIV, my belief as to whether I'm a BIV must track the truth throughout the worlds in which perception plays a substantial role in producing my belief. This means that I can know that I'm not a BIV even though my belief does not track the truth to any BIV worlds. Instead, my belief as to whether I'm a BIV must track the truth from the actual world to the farthest world in which it is true both that I'm not a BIV and that perception plays a substantial role in producing my belief. (For brevity's sake, I will call the latter world *the farthest world*.) Now, in the actual world, as well as in other nearby possible worlds, it is true both that I'm not a BIV and that I believe via perception that I'm not. But what do I believe in the farthest world? I must admit that I don't know exactly how to describe the farthest world, yet it seems that I believe there that I'm not a BIV. Even the farthest world, since it is a world in which perceptual belief-forming methods are available to me, is presumably a world in which I can see my hands (or, at the very least, my nose) when they are before my opened and properly-functioning eyes. Furthermore, it is true in the farthest world that I'm not a BIV. So the counterfactual in condition (4) is true – throughout the worlds in which I'm not a BIV and in which perception substantially produces my belief as to whether I'm a BIV, I believe that I'm not a BIV. So my belief that I'm not a BIV sometimes satisfies conditions (1)-(4) and I can thus know that I'm not a BIV.

But this raises an important question, one that my response must answer if it is to be adequate – why is it sometimes plausible to suppose that I *don't* know that I'm not a BIV? Here's the answer. My belief that I'm not a BIV is alleged to be an *insensitive* one. That is, it is alleged that I would believe that I'm not a BIV even if I were a BIV. And as both Nozick and DeRose suggest, we tend to think that our beliefs do not count as knowledge when we think that they are insensitive. Thus, my belief's insensitivity can explain why we think that it doesn't count as knowledge. But my claim has been that my insensitive belief that I'm not a BIV *does* count as knowledge. So I need to say more in order to explain how we sometimes mistakenly take an insensitive belief that counts as knowledge for one that doesn't. DeRose himself provides two helpful suggestions that we can adapt and use for our own purposes here.

First, let us grant for the moment that my belief that I'm not a BIV is insensitive. This explains why it is sometimes plausible to suppose that I don't know that I'm not a BIV. But if my belief that I'm not a BIV is in fact insensitive, then we are once again faced with the problem of explaining why I *know* that I'm not a BIV. To provide this explanation, we can utilize a suggestion made by DeRose. He claims that I insensitively believe that some skeptical hypotheses are false. For example, I insensitively believe that I'm not a BIV, I insensitively believe that I don't falsely believe that I have hands, and I insensitively believe that I'm not an intelligent dog who is always incorrectly thinking that I have hands.²⁷ But only the first of these insensitive beliefs seems to fail to amount to knowledge. Hence, it is the only one of the three that is effective when used as a premise in an argument like AI. Why are the last two ineffective? DeRose claims that they are ineffective because (1) premises in a skeptical argument are effective when and only when they involve skeptical hypotheses that 'explain how we might come to believe something despite its being

²⁷ See DeRose [8], pp. 196-7.

false',²⁸ and (2) the skeptical hypotheses involved in the last two beliefs do not provide that explanation. Thus, the last two beliefs, despite the fact that they are insensitive, do not seem to fail to amount to knowledge and thus are not effective when used as premises in skeptical arguments.

But how are we to understand effectiveness? DeRose suggests that certain beliefs are effective as premises in skeptical arguments when the skeptical hypotheses they involve are 'persuasive',²⁹ when they pack some punch,³⁰ or when they are 'convincing'.³¹ We might say, then, that *persuasive* skeptical hypotheses are those that explain how we might come to believe something despite its being false. We should note, however, that skeptical hypotheses can be *persuasive* (in this sense) even if they are not *epistemically potent*, that is, even if they fail to strip us of our knowledge. It seems that my knowledge need not be threatened by a skeptical hypothesis simply because that hypothesis explains how I might come to hold certain false beliefs. For example, my visual perceptual knowledge that there is a book on the desk is not threatened by the skeptical hypothesis that I am a large bat who often uses a sophisticated kind of sonar in coming mistakenly to believe that there is a book on the desk. This skeptical hypothesis does not threaten my knowledge even though I insensitively believe that the hypothesis is false, and even though the hypothesis explains how I might come falsely to believe that there is a book on the desk.

What, then, are epistemically potent skeptical hypotheses? If persuasive skeptical hypotheses are those that explain how we might come to hold false beliefs, then in keeping with Nozick's revised analysis of knowledge, I suggest that epistemically potent skeptical hypotheses are those that explain how we might come to hold false beliefs *by using the belief-forming methods that we actually use*. This is, as I have just suggested, consistent with Nozick's revised analysis of

²⁸ DeRose [8], p. 197.

²⁹ DeRose [8], p. 196.

³⁰ See DeRose [8], p. 196.

³¹ DeRose [8], p. 197.

knowledge. Given that analysis, alternative scenarios cannot suggest that I don't know that p unless those scenarios are ones in which I use the belief-forming methods that I actually use. No alternative scenario in which a grandmother mistakenly believes on the basis of *testimony* that her grandson is well can suggest that she doesn't actually know on the basis of *perception* that he is well. So a skeptical hypothesis is epistemically potent, as opposed to merely persuasive, only if it explains how we might come to hold false beliefs by using the belief-forming methods that we actually use.

Given the distinction between persuasive and epistemically potent skeptical hypotheses, we should now note two things. (1) The skeptical hypothesis that I am a BIV is *not* epistemically potent. Since BIVs' belief-forming methods are different from our own, the BIV hypothesis does not explain how I might use the belief-forming methods that I actually use in coming to hold certain false beliefs. Thus, the BIV hypothesis does not have the power to strip us of our knowledge. This means, among other things, that my insensitive belief that I'm not a BIV remains unaffected by the BIV hypothesis. Even though I insensitively believe that I'm not a BIV, the BIV hypothesis does not prevent me from knowing that I'm not a BIV. (2) The BIV skeptical hypothesis can nevertheless be persuasive, for it does explain how I might come to hold certain false beliefs. If I were a BIV, and if I used the belief-forming methods that BIVs use, then it seems that I would falsely believe that I wasn't a BIV. Yet while this does not mean that I don't know that I'm not a BIV, it does explain why it is sometimes plausible to suppose that I don't. Since the BIV skeptical hypothesis explains how I might come to hold certain false beliefs, it can be persuasive. And because it is persuasive, it can create the appearance that my insensitive belief that I'm not a BIV fails to amount to knowledge, which in turn makes my insensitive belief effective as a premise in skeptical

arguments. So the persuasiveness of the BIV hypothesis explains why it is sometimes plausible – especially when I am confronted with that hypothesis – to suppose that I don't know that I'm not a BIV. Nevertheless, all of this is perfectly consistent with saying that we do in fact know that we're not BIVs.

On this picture, then, a skeptical puzzle arises at least in part because the BIV hypothesis is persuasive, in the sense that it explains how we might come to hold certain beliefs despite their being false. And when we put the persuasiveness of the BIV hypothesis together with the apparent truth of AI's second premise and our tendency to provide an affirmative epistemic assessment of my belief that I have hands, we end up with a puzzle. The solution to this skeptical puzzle lies in noticing that even though the BIV hypothesis is *persuasive* and that it is therefore sometimes plausible to suppose that I don't know that I'm not a BIV, the BIV skeptical hypothesis is not epistemically potent and hence does not strip me of the knowledge that I'm not a BIV.

Second, we might dispute the claim that my belief that I'm not a BIV is *insensitive* and claim that it is in fact *sensitive*, that is, that it is *not* the case that I would believe that I'm not a BIV even if I were a BIV. Its sensitivity would then explain why I can know that I'm not a BIV, and our sometimes mistakenly believing it to be *insensitive* would explain why it is sometimes plausible to say that I *don't* know that I'm not a BIV. The proposal now on the table is that S's belief that P is insensitive if S would believe that P even if P were false. But DeRose has another idea.³² Perhaps we should think of insensitivity not in terms of Nozick's initial conditions on knowledge, but in terms of Nozick's revised conditions on knowledge, which take into account methods of belief formation. According to this new proposal, then, S's belief that P is insensitive if S would believe *via method M* that P even if P were false, where M is the method that S actually uses in forming her

³² See DeRose [8], p. 196.

belief. Furthermore, on this new proposal, my belief that I'm not a BIV is *sensitive*. Earlier, I argued that BIVs do not use the same methods of belief formation that we use. Thus, even if I were a BIV, it is *not* the case that I would believe via method M that I'm not a BIV.

This explains why my belief that I'm not a BIV is sensitive and hence how I can know that I'm not a BIV. But why on this new proposal is it sometimes plausible to suppose that I *don't* know that I'm not a BIV? Typically, we fail to realize that BIVs employ methods of belief formation that are different from our own. The experiences of BIVs are, we suppose, exactly similar to ours, and the beliefs that BIVs form on the basis of those experiences seem to be the very same beliefs that we form on the basis of our own experiences. This can lead us to think, as it has led Nozick and others to think, that BIVs use the same methods that we use. But, as I have argued, this is a mistake. Since things like retinas and tympanic membranes help to make our belief-forming methods what they are, our methods are different from those of BIVs. Yet if we mistakenly believe that we share certain methods with BIVs, and if we also think that an adequate definition of insensitivity should take into account methods of belief formation, then we will likely mistakenly believe that my belief that I'm not a BIV is *insensitive*, for we will likely mistakenly think that if I were a BIV, I would believe *via the method that I actually use* that I'm not a BIV. This explains why it is sometimes plausible, although mistaken, to suppose that I don't know that I'm not a BIV.

On this picture, the skeptical puzzle arises at least in part because we mistakenly believe that we share certain belief-forming methods with BIVs. This leads us mistakenly to think that my belief that I'm not a BIV is insensitive and then mistakenly to judge that I don't know that I'm not a BIV. And when we put this together with the apparent truth of AI's second premise and our tendency to provide an affirmative epistemic assessment of my belief that I have hands, we end up with a

puzzle. Here, the solution to the puzzle lies in noticing two things. First, we must notice that we sometimes mistakenly think that we share certain belief-forming methods with BIVs, and that it is therefore sometimes plausible to suppose that I don't know that I'm not a BIV. But second, we must notice that we do not in fact share those methods with BIVs, and hence that my belief that I'm not a BIV can count as knowledge since it is not insensitive.

We have now seen that there are two available proposals about sensitivity and insensitivity. On one proposal, my belief that I'm not a BIV is sensitive, while on the other it is insensitive. Yet no matter which proposal we adopt, and hence no matter whether my belief that I'm not a BIV is sensitive or insensitive, we can adequately explain why it is sometimes plausible to suppose that I don't know that I'm not a BIV. Furthermore, we have explained both how I can know that I have hands and how I can know that I'm not a BIV. So far, then, we have provided two of the three explanations that our Moorean response must provide.

I now turn to the third and final explanation – why is AI's second premise true? It might seem problematic to say within Nozick's framework that it is true, for Nozick *denies* that it is. Nevertheless, DeRose, who adopts and employs parts of Nozick's account, argues that AI's second premise is true. DeRose claims that it is a 'comparative fact that I'm in at least as strong an epistemic position with respect to [the proposition that I'm not a BIV] as I'm in with respect to [the proposition that I have hands]'.³³ This comparative fact 'result[s] in *If I don't know that [I'm not a BIV], then I don't know that [I have hands]* being true regardless of how high or low the standards for knowledge are set'.³⁴ This suggests that, for DeRose, AI's second premise is true even when the standards are always low, and hence even for the purposes of providing a Moorean response to

³³ DeRose [8], p. 203.

³⁴ DeRose [8], p. 203.

skepticism. So we may follow DeRose here in endorsing the truth of AI's second premise within Nozick's framework.

But some contextualists who deny AI's second premise are not motivated by Nozick's considerations. Predominantly, these contextualists are motivated by the relevant alternatives theory of knowledge. Mark Heller claims that we use 'different worlds as relevant alternatives when considering whether [I know that I'm not a BIV] from those used when considering whether [I know that I have hands]'.³⁵ He also suggests that while BIV worlds are *always* relevant to my knowing that I'm not a BIV, they are only *sometimes* relevant to my knowing that I have hands. Those cases in which BIV worlds are not relevant to my knowing that I have hands are cases that demonstrate the falsity of AI's second premise. On our solution, however, the only worlds that are ever epistemically relevant are worlds in which my beliefs are produced by the same methods that produce them in the actual world. Thus, BIV worlds are never relevant to my knowing that I have hands, and they are also never relevant to my knowing that I'm not a BIV. As we've seen, this contributes to my being able to know across contexts both that I have hands and that I'm not a BIV, and thus to our being able to maintain the truth of AI's second premise.

We now have an adequate Moorean response to brain-in-a-vat skepticism, one whose explanations of our epistemic judgements are just as good as those provided by contextualism. Our Moorean response also preserves our ordinary claims to know things about the external world, and it does so without claiming that the skeptic can raise the standards for knowledge and without committing us to disagreeable skeptical conclusions. Given all of this, I conclude that we should favor the Moorean response to skepticism over DeRose's contextualist response.

³⁵ Heller [17], p. 197.

IV. Concluding remarks

At this point, the obvious and natural concern is that the skeptic can build her case without the BIV, whose belief-forming methods are different from my own. Perhaps there are skeptical possibilities that involve creatures whose belief-forming methods are *identical* to mine, and these possibilities might stand a better chance against the kind of Moorean response that we have provided. We certainly should deny neither that there might be such possibilities nor that someone concerned to toe the skeptical line might employ them. But whether or not there are such possibilities, they are typically not employed by skeptics or by those who formulate arguments on the skeptic's behalf. Today, skeptical arguments are usually constructed on a foundation provided by the BIV skeptical hypothesis. Perhaps this indicates that we tend to think that BIV arguments are the toughest skeptical arguments around, or at least that the skeptic needs no tougher argument. In fact, some seem to think that BIV skeptical arguments are so tough that we should give in to them, at least to some extent, and admit that they should be considered either wholly convincing or convincing in certain contexts. So even if the skeptic has other trenches from which she might fight, my Moorean response to arguments like AI represents an advance. For it adequately responds to what some think of as the most troublesome skeptical argument.

Moreover, I suspect that in taking the ground that we have taken, we have discovered some useful strategies and tools that will allow us to advance even further. In particular, we have seen that the belief-forming methods that we actually use play a prominent role in our Moorean response to brain-in-a-vat skepticism. And since the responses that we provide to other skeptical arguments need not be grounded (only) in the claim that other creatures' belief-forming methods are different from our own, other facts about our belief-forming methods might provide the keys to responding to those other skeptical worries. Our responses might be grounded, for example, in whatever it is that

is characteristic of the *beliefs* that are produced by those methods. Or they might be grounded in something that we discover about the *nature* of our belief-forming methods. For instance, they might be grounded in whatever it is that is characteristic of the *experiences* that constitute parts of, or stages in, certain belief-forming methods. In highlighting the fact that our belief-forming methods can be important for responding to the skeptic, our Moorean response to BIV skeptical arguments indicates that investigations into those methods can suggest strategies and tools that will help us to combat other skeptical arguments. And the effectiveness of our response to brain-in-a-vat skepticism might even scare these reserve arguments out of hiding. In any event, however, we must wait to fight these other battles. But we can in the meantime celebrate our victory over one of the skeptic's strongest citadels.³⁶

University of Nebraska-Lincoln

³⁶ For reading and providing invaluable comments on earlier versions of this paper, I am grateful to Keith DeRose, to Ram Neta, and most especially to Albert Casullo. I would also like to thank Robert Audi, Bryan Belknap, Mark Cullison, Joseph Mendola, Peter Murphy, Mark van Roojen, an audience at Harvard University, and two anonymous referees for this Journal for their helpful comments.

REFERENCES

1. Stewart Cohen, 'Knowledge, Context, and Social Standards', *Synthese* 73 (1987), p. 3-26.
2. Stewart Cohen, 'How to be a Fallibilist', *Philosophical Perspectives* 2, *Epistemology* (1988), pp. 91-123.
3. Stewart Cohen, 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* 13 (1999), pp. 57-89.
4. Stewart Cohen, 'Contextualism and Skepticism', *Philosophical Issues* 10 (2000), pp. 94-107.
5. Keith DeRose, 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* 52 (1992), pp. 913-929.
6. Keith DeRose, 'Contextualism: An Explanation and Defense', in John Greco and Ernest Sosa (eds.), *The Blackwell Guide to Epistemology* (Oxford: Blackwell, 1999), pp. 187-205.
7. Keith DeRose, 'Now You Know It, Now You Don't', in *Proceedings of the Twentieth World Congress of Philosophy: Volume V, Epistemology* (Bowling Green, OH: Philosophy Documentation Center, 1999), pp. 1-16.
8. Keith DeRose, 'Solving the Skeptical Problem', in Keith DeRose and Ted A. Warfield (eds.), *Skepticism: A Contemporary Reader* (New York: Oxford University Press, 1999), pp. 183-219.
9. Fred Dretske, 'Epistemic Operators', *Journal of Philosophy* 67 (1970), pp. 1007-1023.
10. Fred Dretske, 'The Pragmatic Dimension of Knowledge', *Philosophical Studies* 40 (1981), pp. 363-378.

11. Alvin I. Goldman, 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73 (1976), pp. 771-791.
12. Alvin I. Goldman, *Epistemology and Cognition* (Cambridge, MA: Harvard University Press, 1986).
13. Alvin I. Goldman, 'Epistemic Folkways and Scientific Epistemology', in *Liaisons: Philosophy Meets the Cognitive and Social Sciences* (Cambridge, MA: MIT Press, 1992), pp. 155-175.
14. Alvin I. Goldman, 'What is Justified Belief?', in Hilary Kornblith (ed.), *Naturalizing Epistemology*, 2nd ed. (Cambridge, MA: MIT Press, 1994), pp. 105-130.
15. John Greco, *Putting Skeptics in Their Place* (Cambridge: Cambridge University Press, 2000).
16. Mark Heller, 'Relevant Alternatives', *Philosophical Studies* 55 (1989), pp. 23-40.
17. Mark Heller, 'Relevant Alternatives and Closure', *Australasian Journal of Philosophy* 77 (1999), pp. 196-208.
18. Mark Heller, 'The Proper Role for Contextualism in an Anti-Luck Epistemology', *Philosophical Perspectives* 13 (1999), pp. 115-129.
19. David Lewis, 'Scorekeeping in a Language Game', *Journal of Philosophical Logic* 8 (1979), pp. 339-359.
20. David Lewis, 'Veridical Hallucination and Prosthetic Vision', *Australasian Journal of Philosophy* 58 (1980), pp. 239-249.
21. David Lewis, 'Elusive Knowledge', *Australasian Journal of Philosophy* 74 (1996), pp. 549-567.

22. Brian P. McLaughlin, 'Lewis on What Distinguishes Perception from Hallucination', in Kathleen Akins (ed.), *Perception* (New York: Oxford University Press, 1996), pp. 198-231.
23. Robert Nozick, *Philosophical Explanations* (Cambridge, MA: Harvard University Press, 1981).
24. Alvin Plantinga, *Warrant and Proper Function* (New York: Oxford University Press, 1993).
25. Steven Rieber, 'Skepticism and Contrastive Explanation', *Noûs* 32 (1998), pp. 189-204.
26. Ernest Sosa, *Knowledge in Perspective* (Cambridge: Cambridge University Press, 1991).