

## DEFENDING A SENSITIVE NEO-MOOREAN INVARIANTISM

Tim Black

California State University, Northridge

In *New Waves in Epistemology*, edited by Vincent F. Hendricks and Duncan Pritchard (Palgrave Macmillan, 2008): 8-27.

I defend a sensitive neo-Moorean invariantism, an epistemological account with the following characteristic features: (a) it reserves a place for a sensitivity condition on knowledge, according to which, very roughly, S's belief that  $p$  counts as knowledge only if S wouldn't believe that  $p$  if  $p$  were false; (b) it maintains that the standards for knowledge are comparatively low; and (c) it maintains that the standards for knowledge are invariant (i.e., that they vary neither with the linguistic context of the subject of knowledge nor with the linguistic context of the attributor of knowledge). I argue that this sort of account allows us to respond adequately to some difficult puzzles in epistemology, puzzles such as skeptical puzzles, as well as puzzles that inspire epistemological contextualism. I also maintain that by utilizing what Keith DeRose calls a warranted assertibility maneuver, sensitive neo-Moorean invariantism can account for our epistemic judgments in each of these puzzle cases.

My aim here is to provide a sketch of an epistemological account that we might euphoniously label a *sensitive neo-Moorean invariantism*. Such an account has the following characteristic features: (a) it reserves a place for a sensitivity condition on knowledge, according to which, very roughly, S's belief that  $p$  counts as knowledge only if S wouldn't believe that  $p$  if  $p$  were false; (b) it maintains that the standards for knowledge are comparatively low; and (c) it maintains that the standards for knowledge are invariant (i.e., that they vary neither with the linguistic context of the subject of knowledge nor with the linguistic context of the attributor of knowledge). My sensitive neo-Moorean invariantism stands apart from fashionable epistemological accounts in at least two ways—it unfashionably makes a place for a condition on knowledge that has its roots in a particularly unfashionable condition, namely, the sensitivity condition, and it unfashionably maintains that the standards for knowledge are invariant across contexts. I therefore do not expect my account to be especially, uh, fashionable.

## 1. Sensitivity? You're Kidding, Right?

But why sensitivity? As Keith DeRose reminds us, sensitivity retains some measure of plausibility even in the face of serious and stubborn objections. Moreover, as I explain below, there are compelling considerations in favor of sensitivity, considerations that give us reason to explore the prospects for providing an account that involves sensitivity.

When we inquire into the nature of knowledge, it seems that if S is to know something about her environment—say, that  $p$ —she must respond appropriately to it, which means at least that she must come to hold the true belief that  $p$ . Yet this cannot be the end of the story, for if S's belief that  $p$  is luckily or coincidentally acquired, it doesn't count as knowledge, not even if it's true and suitably formed. If S is to know that  $p$ , she must do more than respond appropriately to her environment by suitably forming the true belief that  $p$ . It seems that she should also respond appropriately in environments in which it is *not* the case that  $p$ . Robert Nozick puts the point like this:

A belief that  $p$  is knowledge that  $p$  only if it somehow varies with the truth of  $p$ .

... An account that says nothing about what is believed in any situation when  $p$  is false cannot give us any mode of varying with the fact.<sup>1</sup>

But since S's present environment is, *ex hypothesi*, an environment in which  $p$  is true, this additional demand must take the form of a contrary-to-fact conditional: *If  $p$  had not been true, S would have responded differently to her environment and, in particular, would not have believed that  $p$ .*

As this suggests, it's true not only that these conditions—truth and belief—must be satisfied if one is to know, but also that the conditions themselves be related. In particular, S's

---

<sup>1</sup> Nozick (1981), p. 208.

knowing that  $p$  calls for an intimate connection between whether  $p$  is true and whether S believes that  $p$ . Any adequate epistemological account should have something to say about the relationship between truth and belief, and a sensitivity account does so: it straightforwardly captures the idea that if one is to know, one's beliefs must vary with, or conform to, the truth.

Other theories have trouble capturing this idea. The reliabilist, for example, who maintains that S must respond appropriately to her environment if she is to know something about it, is concerned in the first place with the reliability rating of the process that generates S's belief. And given that S does in fact believe that  $p$ , the reliabilist will have much to say about whether her belief is justified: S's belief is justified when it is generated by a reliable belief-forming process; and if the belief is true, the reliabilist might be glad to count it as knowledge. Yet suppose that  $p$  were *not* true. The proponent of sensitivity maintains, as we've seen, that S wouldn't believe that  $p$  if  $p$  weren't true. The reliabilist, however, has surprisingly little to say here. How, according to the reliabilist, would S perform if  $p$  were false? One might think that the reliabilist will maintain either that S would believe that  $p$  is false or that S would withhold belief as to whether  $p$ . But the reliabilist is simply silent on this issue, and so it seems that the proponent of sensitivity can better account for the connection between truth and belief as well as for S's doxastic performance in situations in which  $p$  is false.<sup>2</sup>

Sensitivity also helps us deal with cases that count against reliabilism. Suppose that Henry is driving through a region populated by *papier-mâché* barn facsimiles. Even if Henry

---

<sup>2</sup> Still, rather than saying what S *would* believe, the reliabilist might maintain that S *should* believe that  $p$  is false or that S *should* withhold belief as to whether  $p$ . And although the reliabilist can provide this sort of guidance, her capacity to do so cannot help her to account for the connection that presently concerns us, viz., the connection between belief and *truth*. In fact, it seems that the reliabilist, who suggests that our doxastic performances should be directed by considerations of reliability rather than by considerations of truth, will be hard pressed to account for the connection between belief and truth. Thus, for instance, there will almost certainly be situations in which, guided by considerations of reliability, the reliabilist would direct S to believe that  $p$  even if  $p$  were not true.

spies a bona fide barn—even if he believes, justifiedly and on the basis of a reliable belief-forming process, that this is a barn—it seems that Henry does not *know* that it’s a barn.

Sensitivity can explain why Henry fails to know. For if this were *not* a barn, Henry would nevertheless have believed that it *is* a barn—Henry doesn’t know that this is a barn because he insensitively believes that it’s a barn.

We can strengthen the case for sensitivity by appealing to additional cases. A sensitivity account handles, for example, Gettier’s original cases against the justified-true-belief account of knowledge.<sup>3</sup> Suppose that Smith is justified in believing that

(f) Jones own a Ford.

From this, Smith justifiedly infers that

(h) Either Jones owns a Ford, or Brown is in Barcelona.

Since Smith is justified in believing that (f), and since the inference from (f) to (h) is perfectly legitimate, Smith is justified in believing that (h). Furthermore, suppose that (h) is true. Smith therefore has a justified true belief that (h). But Smith doesn’t *know* that (h). For “Jones does *not* own a Ford, [and] ... by the sheerest coincidence, and entirely unknown to Smith,”<sup>4</sup> Brown is indeed in Barcelona. Unlike the justified-true-belief account, however, a sensitivity account can explain why Smith fails to know that (h). For even though Smith’s belief that (h) is true, he would believe that (h) even if it were false. Since Smith’s belief doesn’t count as knowledge because it is insensitive and it doesn’t suitably vary with the truth.

The literature that sprang up in the wake of Gettier’s paper provides a nice array of additional test cases. For example, in Brian Skyrms’ Sure-Fire match case, which concerns

---

<sup>3</sup> See Gettier (1963).

<sup>4</sup> Gettier (1963) (quoted from Roth and Galis (1970), p. 38).

inductive knowledge and knowledge of the future, a pyromaniac justifiedly believes that the next Sure-Fire match he strikes will ignite.<sup>5</sup> However, the match is defective—it ignites, but only because it was hit by a coincidental burst of Q-radiation. The pyromaniac has a justified true belief that the match will ignite, but he doesn't *know* that it will. Here again, sensitivity handles the case: if the match hadn't ignited—if there had been no coincidental burst of Q-radiation—the pyromaniac would nevertheless have believed that the next match he strikes will ignite. The pyromaniac's belief fails to count as knowledge because it is insensitive.

These cases suggest that we should at least explore the prospects for providing an account of knowledge that involves sensitivity. Additional reasons are supplied by the theoretical considerations at the beginning of this section: we should build our account of knowledge around sensitivity in order to ensure that the only beliefs that count as knowledge are those that suitably vary with the truth. These considerations give us reason to pursue an account of knowledge that involves sensitivity.<sup>6</sup>

## 2. Skeptical Puzzles

I have claimed that according to my account, the standards for knowledge are both comparatively low and invariant. To bring this out, while also bringing out my account's ability

---

<sup>5</sup> See Skyrms (1967).

<sup>6</sup> Peter Murphy and I (in Black and Murphy (ms.)) explore other reasons to pursue an account of knowledge that involves sensitivity. In particular, we propose a sensitivity account of knowledge that stands up against the principal objection to such accounts, namely, that they are incompatible with extremely plausible instances of closure (according to which, roughly, if S knows that *p* and that *p* entails *q*, then S knows that *q*). We consider four cases in which Nozick's account forces us absurdly to deny closure, and we argue that our proposal is in fact compatible with such instances of closure. We argue that our proposal, unlike Nozick's, handles each case, and that it handles each in an instructive way, teaching lessons not only about our proposal in particular but also about any proper construal of the conditions on knowledge. We also argue that our proposal helps us both to clarify and to focus the debate over skepticism.

to respond to some difficult puzzles, I want now to consider a puzzle of which DeRose—and several others—has reminded us. Consider the following skeptical argument:

ARGUMENT FROM IGNORANCE (AI)

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

Therefore,

- (3) I don't know that I have hands.<sup>7</sup>

AI presents a puzzle: From its two plausible premises, we may validly deduce a claim whose *negation* is plausible. Yet in spite of the fact that they are independently plausible, (1), (2), and the negation of (3) are jointly inconsistent. It seems, therefore, that we must give up one of those claims. But which one should we give up, and why? This is *the skeptical puzzle*.

My sensitive neo-Moorean invariantism allows us to provide a response to brain-in-a-vat skepticism. To show that this is the case, I must provide two 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 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.<sup>8</sup>

---

<sup>7</sup> See DeRose (1999b), p. 183.

<sup>8</sup> I must also maintain that AI's second premise is true. It might seem that this will be difficult to do while espousing a sensitivity account of knowledge. For Nozick, who espouses such an account, *denies* that AI's second premise is true. Still, 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]" (DeRose (1999b), p. 203). 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" (DeRose (1999b), p. 203). This suggests that 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 skepticism. Thus, for the purposes of providing our Moorean response to skepticism, we may endorse the truth of AI's second premise within Nozick's framework.

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 the conditions on knowledge, which include holding a true sensitive belief. 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 is sensitive, which allows me to know that I have hands.

Second, my belief that I'm not a BIV can also count as knowledge. For one thing, it's *true* that I'm not a BIV. Moreover, I *believe* that I'm not a BIV. And, as I'll now argue, my belief that I'm not a BIV can be *sensitive*. According to the proposal that's now on the table, S's belief that *p* is sensitive if and only if S wouldn't believe that *p* if *p* were false. But DeRose has another idea:<sup>9</sup> Perhaps we should think of sensitivity not in terms of the current proposal, but in terms of Nozick's revised sensitivity condition, which takes into account methods of belief formation.<sup>10</sup> According to Nozick's revision, S's belief that *p* is sensitive if and only if S wouldn't believe *via method M* that *p* if *p* were false, where *M* is the method that S actually uses in forming her belief. The difference between the original proposal and Nozick's revision is this: In the revision, to determine whether S's belief that *p* is sensitive, we consider the closest possible world in which *p* is false—call this world *w*—and then determine not, as in the original proposal, whether S believes in *w* that *p*, but whether S believes *via M* that *p*. Her belief is

---

<sup>9</sup> See DeRose (1999b), p. 196.

<sup>10</sup> DeRose (1999b), p. 196. See also Nozick (1981), p. 179.

sensitive if *either* (a) she does not believe in  $w$  that  $p$  or (b) in  $w$ , she believes that  $p$  via some method other than  $M$ .

Let's grant that in the closest worlds in which I'm a BIV, I believe that I'm not a BIV. Nozick's revised condition nevertheless allows for the possibility that my belief is sensitive. To see whether my belief *is* sensitive, we need to determine how I come to believe in the closest BIV worlds that I'm not a BIV. If I form that belief by using some method other than the method that I actually use in forming that belief, then my belief that I'm not a BIV is sensitive. I want now to argue that *BIVs do not use our belief-forming methods*. Thus, even if I were a BIV, I would not believe via  $M$  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.

To see that BIVs do not use our belief-forming methods, we must first circumscribe the sphere of worlds throughout which either perception or an inference from perceptual beliefs produces my belief as to whether I'm a BIV. For, as things are, 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. Once we've circumscribed the appropriate sphere of worlds, we'll see that there are no BIV worlds in this sphere, for in such worlds 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 use a sophisticated artificial process to stimulate my brain electrochemically. This method is *not* the method of perception.

Of course, this last 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.<sup>11</sup>

But Nozick’s conclusion—that I and the BIV use the same belief-forming method—is 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.<sup>12</sup>

There is a problem with this analysis: It counts as visual perception some methods that seem not to be visual perception.

Suppose that Ray has been blind from birth. In spite of this, he one day 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 method—and we may suppose that his experience, which is the only

---

<sup>11</sup> Nozick (1981), pp. 184-185.

<sup>12</sup> 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.

experience of its kind he has ever had, is the same ‘from the inside’ as a visual experience. (Let’s say that it’s 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 method counts as visual perception. But this seems to be a mistake; ordinarily, we would not say that Ray’s method counts as visual perception. We conclude only that Ray has had an experience of a certain sort, not that his method counts as visual perception. Nozick’s account of visual perception is therefore unsatisfactory. If a method is to count as visual perception, it must not only be experientially the same, the same ‘from the inside’, as visual perception, but it must also meet some additional conditions. And these additional conditions must be ‘outside’ conditions.

This allows us to say, then, as 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.<sup>13</sup>

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 don’t count as visual perception, for example, from Ray’s 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 method, we do not count that method as visual perception. Furthermore, we should also

---

<sup>13</sup> McLaughlin (1996), 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 (1980).

characterize other perceptual belief-forming methods ‘from the outside’ if we want to distinguish them from methods that don’t 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.<sup>14</sup>

Note now that 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 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. So, 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. This means that my belief that I’m not a BIV is sensitive, and that it can count as knowledge.

### **3. Persuasiveness and Epistemic Potency**

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? Well, it’s alleged

---

<sup>14</sup> 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.

that my belief that I'm not a BIV is *insensitive*, and we tend to think that our beliefs don't count as knowledge when they are insensitive.<sup>15</sup> So, according to the allegation, it's sometimes plausible to suppose that I don't *know* that I'm not a BIV because I *insensitively believe* that I'm not a BIV.

This, however, is a mistake: it is *not* the case that I insensitively believe that I'm not a BIV. There are two reasons, as I see things, why we might make the mistake of thinking that I insensitively believe that I'm not a BIV. First, we might fail to recognize that sensitivity involves belief-forming methods, in which case we might conceive of sensitivity as according to the original proposal. This would lead us, of course, to think that I insensitively believe that I'm not a BIV. Second, we might recognize that sensitivity involves belief-forming methods, but fail to recognize that BIVs' methods are not identical to those that I use in coming to believe that I'm not a BIV. This too would lead us to think that my belief is insensitive. Let's begin with the second of these ideas.

We might 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

---

<sup>15</sup> See DeRose (1999), p. 193; and, of course, Nozick (1981).

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 insensitively believe 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 with 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 puzzle lies in realizing two things. First, we must realize that we sometimes mistakenly think that we share certain belief-forming methods with BIVs, and that we can therefore be led to suppose that I insensitively believe that I'm not a BIV. But second, we must realize that we do not in fact share those methods with BIVs, and hence that my belief that I'm not a BIV is *not* insensitive.

Let's move on to the second reason why we might mistakenly think that I insensitively believe that I'm not a BIV: We might make such a mistake because we fail to conceive of sensitivity in accordance with Nozick's revision, which incorporates belief-forming methods, and conceive of it instead in accordance with our original proposal. But how does this lead us mistakenly to think that my belief is insensitive? To answer this question, it will help to consider some thoughts from DeRose, who thinks that I insensitively believe that I'm not a BIV.<sup>16</sup>

DeRose claims that I insensitively believe that some skeptical hypotheses are false. He claims, for example, that I insensitively believe that I'm not a BIV, that I insensitively believe

---

<sup>16</sup> See DeRose (1999), pp. 193, 196-197.

that I don't falsely believe that I have hands, and that I insensitively believe that I'm not an intelligent dog who is always incorrectly thinking that I have hands.<sup>17</sup> But only the first of these beliefs seems 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 false",<sup>18</sup> and (2) the skeptical hypotheses involved in the last two beliefs don't provide that explanation.

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".<sup>19</sup> Moreover, he seems to suggest in (1) that a skeptical hypothesis, H, is persuasive only if it explains how in scenarios in which H is true, I would nevertheless have reason to believe that H is false (or to hold a belief that is incompatible with the belief that H). Thus, the skeptical hypothesis that I have no hands (i.e. that I falsely believe that I have hands) is *not* persuasive, for it fails to explain how in scenarios in which I have no hands, I would nevertheless have reason to believe that I *do* have hands (i.e. that I don't falsely believe that I have hands). If we go on to characterize sensitivity in terms of persuasiveness, we can say that the last two of DeRose's three hypotheses are ineffective—and therefore unpersuasive—because in spite of appearances, those hypotheses are in fact sensitive.

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 it is

---

<sup>17</sup> See DeRose (1999b), pp. 196-197.

<sup>18</sup> DeRose (1999b), p. 197.

<sup>19</sup> DeRose (1999b), p. 196.

persuasive. So, for example, my visual perceptual knowledge that there is a book on the desk is not threatened by the skeptical hypothesis that

(BAT) I am a large bat who often uses a sophisticated kind of sonar in forming beliefs about the external world but who mistakenly believes on the basis of its sonar that there is a book on the desk.

(BAT) does not threaten my visual perceptual knowledge that there is a book on the desk even though (BAT) explains how in scenarios in which (BAT) is true, I would have a reason to believe that it's false.

What, then, are epistemically potent skeptical hypotheses? Let's begin to answer this question by supposing that in arguing for skeptical conclusions, we argue that a source of beliefs, *M*, has no power to justify the beliefs—or a belief—that it produces. Consider two beliefs, my belief that *x* and my belief that *y*, where *x* and *y* are incompatible, and where my belief that *x* is formed via *M*. We will show that *M* fails to justify my belief that *x* only if we show that in a (counterfactual) situation in which *y* is true, we still have a reason to believe *via M* that *y* is false. That is to say, we will show that *M* fails to justify my belief that *x* only if we show that even when *x* is false, I would nevertheless have a reason to believe *via M* that *y* is false (i.e., a reason to believe *via M* that something compatible with *x*—and perhaps *x* itself—is true).

This qualification—the *via-M* qualification—must be included in our characterization of epistemic potency. For suppose that *M*—say, visual perception—produces my belief that *x*, while *M\**—say, the testimony of a person, *S*—produces my belief that *y* is false. We learn from Nozick, though, that my visual perceptual knowledge that *x* is not threatened by the fact that in a situation in which *y* is true, I would have a reason to believe *via S's* testimony that *y* is false: 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.<sup>20</sup> To take another example, my knowledge that there is a book on my desk, when I acquire that knowledge by looking directly at the book in optimal lighting conditions and with my properly-functioning eyes, is not threatened by the fact that in a counterfactual situation in which there is no book on my desk, I would have a reason to believe via a colleague's testimony that there *is* a book on my desk. These examples strongly suggest that our characterization of epistemic potency should include the via-M qualification.

Here, then, is that characterization: If a skeptical hypothesis is persuasive only if it explains how in scenarios in which the hypothesis is true, I would nevertheless have a reason to believe that it's false, then in keeping with Nozick's revision, I suggest that

(EP) a skeptical hypothesis, H, is *epistemically potent* (against my belief that *p*) only if it explains how in scenarios in which H is true, I would nevertheless have a reason to believe that H is false *via the method that I actually use in coming to hold the belief that p*.

(EP) is consistent with Nozick's revision. On that revision, 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.

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 in scenarios in which I am a BIV, I would have a reason to believe that I'm not a BIV via the method that I actually use in coming to hold the belief that I have hands. Thus, the

---

<sup>20</sup> See Nozick (1981), p. 179.

BIV hypothesis does not have the power to strip me of my knowledge that I have hands. Moreover, the BIV hypothesis doesn't prevent me from knowing that I'm not a BIV. For, since BIVs' belief-forming methods are different from our own, the BIV hypothesis does not explain how in scenarios in which I am a BIV, I would have a reason to believe that I'm not a BIV via the method that I actually use in coming to hold that belief. If we now go on to characterize sensitivity in terms of epistemic potency, we can say that the BIV hypothesis is ineffective—and therefore unpersuasive—because in spite of appearances, that hypothesis is in fact sensitive. (2)

The BIV skeptical hypothesis can nevertheless be persuasive, for it does explain how in scenarios in which I am a BIV, I would have a reason to believe that I'm not a BIV: If I were a BIV, and if I used the belief-forming methods that BIVs use, I would have a reason to believe that I'm not a BIV. Yet while this doesn't mean that I insensitively believe that I'm not a BIV, it does explain why it's sometimes plausible to suppose that I do. Since the BIV skeptical hypothesis can be persuasive, it can create the appearance that my belief that I'm not a BIV is insensitive, and hence that it fails to amount to knowledge, which can in turn make that belief effective as a premise in skeptical arguments. So the *persuasiveness* of the BIV hypothesis explains, among other things, why we sometimes suppose—especially, perhaps, when we're confronted with that hypothesis—that we *insensitively believe* that we're not BIVs. Nevertheless, all of this is perfectly consistent with saying that my belief that I'm not a BIV is sensitive, and that I do in fact know that I'm not a BIV.

On this picture, then, the skeptical puzzle arises at least in part because the BIV hypothesis is persuasive, in the sense that it explains how in scenarios in which I am a BIV, I would have a reason to believe that I'm not a BIV. This can lead 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. When we put this together with the apparent truth of AI's second premise and with 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 puzzle lies in realizing two things. First, we must realize that the BIV hypothesis is persuasive and that we can therefore be led to suppose that we insensitively believe that we're not BIVs. But second, we must realize that the BIV skeptical hypothesis is *merely* persuasive; it is *not* epistemically potent. And this allows us to maintain that my belief that I'm not a BIV is *not* insensitive, and therefore that I can know that I'm not a BIV.

#### 4. Warranted Assertability

We've now seen that we can employ the notion of persuasiveness in order to explain why it's sometimes plausible to suppose that we insensitively believe that certain skeptical hypotheses are false and consequently why it's sometimes plausible to suppose that we don't know that they're false. We can also, at this point in our meandering dialectic, offer another explanation of the plausibility of certain denials of knowledge, an explanation that appeals to the notion of *warranted assertability*.<sup>21</sup> Through another example, we can see how that notion can help to explain the plausibility of certain denials of knowledge.

Consider Cohen's airport case:

---

<sup>21</sup> In providing this explanation, not only will we bolster our case for invariantism, but we will also establish a connection between persuasiveness and warranted assertability: A skeptical hypothesis, H, is persuasive in those cases in which we're warranted in asserting that S doesn't know that *p*, where H entails  $\sim p$ .

There are other invariantist explanations of the plausibility of certain denials of knowledge. For example, Pritchard (Forthcoming) rejects the contextualist's linguistic intuition that in response to mere conversational changes, we are willing to move from asserting that we don't know to asserting that we do know, or from asserting that we do know to asserting that we don't. He argues that instead of moving from one categorical assertion to another, we move (for example) from making a categorical assertion of knowledge to making no assertion at all (or to making only a qualified assertion). He goes on to argue that this linguistic data supports an invariantist interpretation as well as a Gricean explanation, in terms of the notion of warranted assertability, of the plausibility of certain denials of knowledge. For other explanations, see Bach (Forthcoming) and Hawthorne (2004).

Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, “Yes I know—it does stop in Chicago.” It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, “How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute.” Mary and John agree that Smith doesn’t really *know* that the plane will stop in Chicago. They decide to check with the airline agent.<sup>22</sup>

This case involves two conversational contexts.<sup>23</sup> The first context is one in which Smith says, “Yes I know—it does stop in Chicago.” The second is one in which Mary says, “Smith doesn’t really know that the plane will stop in Chicago.” Both utterances are about Smith’s belief that the plane will stop in Chicago: Smith says that he knows that the plane will stop in Chicago; and Mary says that he doesn’t know. The question that faces us, then, is this: Why is it plausible for Mary to say that Smith doesn’t know that the plane will stop in Chicago?

Before we answer this question, though, I should say that the sensitive neo-Moorean invariantist maintains that Mary says something *false*: Even in her conversational context, Smith knows that the plane will stop in Chicago. The concerns Mary voices when she presents her skeptical hypotheses are not epistemically potent, for neither hypothesis explains how, in the scenarios in which it is true, Smith would have a reason to believe that the relevant hypothesis is

---

<sup>22</sup> Cohen (1999), p.58.

<sup>23</sup> DeRose points out that there are complications with the first context. These complications arise from the fact that the ascriptions in that context are first-person in nature. See DeRose (2005), Section 5.

false *via the method that he uses in coming to believe that the plane will stop in Chicago*. Smith forms the belief that the plane will stop in Chicago via the testimony of his itinerary, but only a *different* method would give him a reason to believe that the itinerary contains no misprint.<sup>24</sup> For the testimony of Smith's itinerary, even though it is the source of his belief that the plane will stop in Chicago, gives him no reason to believe that the itinerary itself is reliable, that is, that the itinerary itself contains no misprint. To have a reason to believe that the itinerary contains no misprint, Smith must consult a different source, a source other than the itinerary itself. In general, it seems, a source,  $s_1$ , can give S no reason to believe that  $s_1$  itself is reliable. In order to have such a reason, S must consult a different, independent (of  $s_1$ ) source,  $s_2$ . So, since the testimony of Smith's itinerary—that is, the method that he uses in coming to believe that the plane will stop in Chicago—gives him no reason to believe that the itinerary contains no misprint, Mary's hypothesis that the itinerary contains a misprint is not epistemically potent. It therefore presents no barrier to Smith's knowing, even in Mary's conversational context, that the plane will stop in Chicago.

Back, then, to the question why it is plausible for Mary and John to suppose that Smith doesn't know that the plane will stop in Chicago. To answer this question, I'll appeal to the notion of warranted assertability. It's common in explaining why certain assertions are warranted in certain contexts to appeal to Grice's conversational maxims, and I'll do so here. I remind the reader in particular of Grice's second maxim of Quantity,

Q2: Do not make your contribution more informative than is required.<sup>25</sup>

---

<sup>24</sup> This line of reasoning, *mutatis mutandis*, applies equally well to the hypothesis that the powers that be could have changed the schedule at the last minute. Since this is the case, I focus for brevity's sake on the hypothesis that the itinerary could contain a misprint.

<sup>25</sup> Grice (1989), p. 26.

I can violate Q2 by providing irrelevant information. Suppose S asks, “Is John taller than Fred?” All else being equal, I violate Q2 when I respond by saying, “Yes. John is taller than Fred, and it’s unseasonably cool in Los Angeles.”<sup>26</sup> Yet this is not the only way to violate Q2: I can violate Q2 by providing information that, even though it’s relevant, exceeds the threshold of informativeness. Call a submaxim of the latter sort a *submaxim of strength*. In the next two paragraphs, I will sketch an explanation—one that appeals to a submaxim of strength—of the plausibility of Mary’s supposing that Smith doesn’t know that the flight stops in Chicago.

Mary and John’s conversation concerns the practical question of whether they’ll be able to meet their business contact in Chicago. Mary also raises the possibilities that the itinerary contains a misprint and that the powers that be have changed the flight schedule at the last minute. Given the stakes and the mentioned error possibilities, what is salient to their conversation is a strong epistemic position, one strong enough to rule out even unlikely errors, such as the possibility that the itinerary contains a misprint.<sup>27</sup> But Smith cannot eliminate these error possibilities, and the fact that they’re salient makes it seem (to Mary and John) that Smith must eliminate them if he is to know that the flight stops in Chicago. Moreover, Mary and John’s assertion ‘Smith doesn’t know that the flight will stop in Chicago’ conveys the information that Smith cannot eliminate all of those salient possibilities. Thus, since he is in fact in no position to eliminate all of those possibilities, Mary and John’s assertion is made in accordance with the appropriate submaxim of strength—it doesn’t exceed the threshold of informativeness. They may

---

<sup>26</sup> This example is similar to one provided by van Rooy (ms.), p. 18.

<sup>27</sup> For more on the distinction between salience and relevance, see Rysiew (2001), p. 490. I mean to claim here that salient alternatives affect only the assertability conditions of knowledge attributions, while only relevant (or epistemically potent) alternatives can affect their truth conditions. In Mary’s conversational context, the alternative that the itinerary contains a misprint is *salient*—and affects only the assertability conditions of certain attributions (of knowledge or of a lack of knowledge)—but it is *not relevant*—and so it doesn’t affect the truth conditions of knowledge attributions.

therefore assert that Smith doesn't know that the flight stops in Chicago.

How does the warrantedness of Mary's assertion create the appearance that it's false that Smith knows that the plane will stop in Chicago? In accordance with the Gricean submaxim of strength, Mary is warranted in asserting in her conversational context that Smith doesn't know that the flight will stop in Chicago. This warranted assertion can generate the false implicature that Smith's belief that the plane will stop in Chicago does *not* meet the standards for knowledge. This false implicature can in turn create the appearance that it's false that Smith knows that the plane will stop in Chicago. In this way, then, the warrantedness of Mary's assertion can lead us mistakenly to suppose that it's false that Smith knows.<sup>28</sup>

## 5. The Adequacy of the Strength-Based Explanation

Is this explanation satisfactory? So far as I know, only DeRose has proposed criteria of adequacy for an explanation of this sort, which involves what he calls a warranted-assertability maneuver, or WAM.<sup>29</sup> He suggests that

- (I) WAMs are successful only if they need to explain away *only* an appearance of falsity—and, in particular, an appearance of falsity that conflicts, in one and the same context, with another appearance of falsity;
- (II) WAMs are successful only if they can provide the needed explanation(s) by appealing only to the generation of a false implicature (or false implicatures); and
- (III) WAMs are successful only if they appeal to *general* rules of conversation in explaining why certain assertions are unwarranted (or warranted).

---

<sup>28</sup> I provide the details of this account in Black (ms.).

<sup>29</sup> See DeRose (1999a), especially pp. 201-203. Jessica Brown, too, is concerned to provide a WAM that can either meet or plausibly resist each of these three conditions; see Brown (2005), pp. 283-285.

Our strength-based WAM meets all three of these conditions. First, it begins with intuitions of falsity that conflict intracontextually, that is, within one and the same context. Recall Mary's conversational context, which is characterized by her presenting certain skeptical hypotheses that, according to the contextualist, put in place unusually high epistemic standards. In Mary's context, it seems false that Smith knows that the plane will stop in Chicago. After all, he can't rule out the salient possibilities that the itinerary contains a misprint and that the powers that be have changed the schedule at the last minute.

But doesn't it also seem false that Smith *doesn't* know that the plane will stop in Chicago? As DeRose himself says, while speaking of radical skeptical contexts, "[m]ost of us feel some ambivalence" here.<sup>30</sup> For we realize, as does Mary, that in coming to believe that the plane will stop in Chicago, Smith relies on the testimony of his itinerary, which is a rather trustworthy source. We also realize that he has no reason in Mary's context to think that his itinerary contains a misprint, for all that's happened in that context is this: Mary *mentions the possibility* that the itinerary contains a misprint. I submit, then, that we should maintain that both of the following hold, simultaneously and intracontextually, in Mary's conversational context: (A) It seems false to say that Smith *knows* that the plane will stop in Chicago, and (B) it seems false to say that he *doesn't* know that the plane will stop in Chicago. The invariantist, then, must explain away the appearance in (A). But since this is an appearance of falsity, the invariantist satisfies DeRose's first condition.

Second, our invariantist WAM explains away this appearance of falsity by appealing only to the generation of false implicatures. The warrantedness of Mary's assertion that Smith doesn't know that the plane will stop in Chicago generates the false implicature that his belief fails to

---

<sup>30</sup> DeRose (1999b), pp. 208-9. (Compare DeRose (1999b), p. 185.)

meet the standards for knowledge, which causes us mistakenly to suppose that it's false that he does know.

Finally, in explaining why Mary's assertion that Smith doesn't know that the plane will stop in Chicago is warranted, our WAM appeals to very general conversational rules, namely, to Grice's maxims of Quantity and to submaxims of strength. So, since our WAM meets all three of DeRose's conditions, both he and those who accept DeRose's criteria of adequacy should count it as satisfactory.

## **6. Conclusion**

We've now seen that a sensitive neo-Moorean invariantism, according to which the standards for knowledge are both invariant and comparatively low, allows us to respond to brain-in-a-vat skepticism. In showing that this is the case, I explained how we can know that we have hands, and how we can know that we're not BIVs. Moreover, I explained how the persuasiveness of the BIV skeptical hypothesis causes it sometimes to be plausible to say that we *don't* know that we're not BIVs. Still, even though the BIV hypothesis is persuasive, it is not epistemically potent and hence does not strip us of the knowledge that we're not BIVs. Finally, I performed a warranted-assertability maneuver in order to explain the plausibility of a denial of knowledge in Cohen's airport case, a case that is often taken to support (or even to establish) epistemological contextualism. This WAM, which meets all the criteria of adequacy for WAMS, is rooted in Grice's second maxim of Quantity and, in particular, in a submaxim of strength. I take it, then, that we have established the plausibility of a sensitive neo-Moorean invariantism. Such an invariantism gives us the tools we need to respond to skeptical challenges, while at the same time it rises to the occasion when called on to provide the explanations that it must provide. On the

strength of these considerations, then, we should give sensitive neo-Moorean invariantism a place among the most attractive epistemological accounts.

#### REFERENCES

Bach, Kent (Forthcoming) 'The Emperor's New 'Knows'', in Gerhard Preyer and Georg Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (Oxford UP).

Black, Tim (ms.) 'A Warranted-Assertability Defense of a Moorean Response to Skepticism' [online at <http://www.csun.edu/~tab2595/WAMs.pdf>].

Black, Tim and Peter Murphy (ms.) 'Explanationist Counterfactualism'.

Brown, Jessica (2005) 'Adapt or Die: the Death of Invariantism?' *The Philosophical Quarterly* 55: 263-285.

DeRose, Keith (1999a) 'Contextualism: An Explanation and Defense', in John Greco and Ernest Sosa (eds.), *The Blackwell Guide to Epistemology* (Blackwell).

DeRose, Keith (1999b) 'Solving the Skeptical Problem', in Keith DeRose and Ted A. Warfield (eds.), *Skepticism: A Contemporary Reader* (Oxford UP).

DeRose, Keith (2005) 'The Ordinary Language Basis for Contextualism and the New

Invariantism' *The Philosophical Quarterly* 55: 172-198.

Gettier, Edmund (1963) 'Is justified true belief knowledge?' *Analysis* 23: 121-123. Reprinted in Roth and Galis (1970), pp. 35-38.

Grice, H. P. (1989) 'Logic and Conversation', in H. P. Grice, *Studies in the Way of Words* (Harvard UP).

Hawthorne, John (2004) *Knowledge and Lotteries* (Oxford UP).

Lewis, David (1980) 'Veridical Hallucination and Prosthetic Vision', *Australasian Journal of Philosophy* 58: 239-249.

McLaughlin, Brian P. (1996) 'Lewis on What Distinguishes Perception from Hallucination', in Kathleen Akins (ed.), *Perception* (Oxford UP).

Nozick, Robert (1981) *Philosophical Explanations* (Harvard UP).

Pritchard, Duncan (Forthcoming) 'Contextualism, Skepticism and Warranted Assertability Manœuvres', in J. Keim-Campbell, M. O'Rourke and H. Silverstein (eds.), *Knowledge and Skepticism*, (MIT Press).

Roth, Michael D. and Leon Galis, eds. (1970) *Knowing: Essays in the Analysis of Knowledge*

(Random House).

Rysiew, Patrick (2001) 'The Context-Sensitivity of Knowledge Attributions', *Noûs* 35: 477-514.

Skyrms, Brian (1967) 'The Explication of 'X Knows That *p*'', *Journal of Philosophy* 64: 373-389. Reprinted in Roth and Galis (1970), pp. 89-111.

van Rooy, R. (ms.) 'Relevance Implicatures' [online at <http://turing.wins.uva.nl/~vanrooy/Implicfinal.pdf>].