

MYLAN ENGEL, Jr.

## WHAT'S WRONG WITH CONTEXTUALISM, AND A NONCONTEXTUALIST RESOLUTION OF THE SKEPTICAL PARADOX

**ABSTRACT.** Skeptics try to persuade us of our ignorance with arguments like the following: 1. I don't know that I am not a handless brain-in-a-vat (BIV). 2. If I don't know that I am not a handless BIV, then I don't know that I have hands. Therefore, 3. I don't know that I have hands. The BIV argument is valid, its premises are intuitively compelling, and yet, its conclusion strikes us as absurd. Something has to go, but what? Contextualists contend that an adequate solution to the skeptical problem must: (i) retain epistemic closure, (ii) explain the intuitive force of skeptical arguments by explaining why their premises initially seem so compelling, and (iii) account for the truth of our commonsense judgment that we do possess lots of ordinary knowledge. Contextualists maintain that the key to such a solution is recognizing that the semantic standards for "knows" vary from context to context such that in skeptical contexts the skeptic's premises are true and so is her conclusion; but in ordinary contexts, her conclusion is false and so is her first premise. Despite its initial attractiveness, the contextualist solution comes at a significant cost, for contextualism has many counterintuitive results. After presenting the contextualist solution, I identify a number of these costs. I then offer a noncontextualist solution that meets the adequacy constraint identified above, while avoiding the costs associated with contextualism. Hence, one of the principal reasons offered for adopting a contextualist theory of knowledge – its supposedly unique ability to adequately resolve the skeptical problem – is undermined.

The *intense* view of these manifold contradictions and imperfections in human reason has so wrought upon me, and heated my brain, that I am ready to reject all belief and reasoning, and can look upon no opinion even as more probable or likely than another. Where am I, or what? ...

Most fortunately it happens, that since reason is incapable of dispelling these clouds, nature herself suffices to that purpose, and cures me of this philosophical melancholy and delirium, either by relaxing this bent of mind, or by some avocation, and lively impression of my senses, which obliterate all these chimeras. I dine, I play a game of back-gammon, I converse, and am merry with my friends; and when after three or four hour's amusement, I wou'd return to these speculations, they appear so cold, and strain'd, and ridiculous, that I cannot find in my heart to enter into them any farther (Hume 1739, p. 268f).

There is an inconvenience which attends all abstruse reasoning, that it may silence, without convincing an antagonist, ... When we leave our closet, and engage in the common affairs of life, its conclusions seem to vanish, like the phantoms of the night on the appearance of the morning (Hume 1739, p. 455).

David Hume



## 1. INTRODUCTION

Like the phantom-free Hume who has left his closet and is basking the light of commonsense, I know a lot. I know that I have hands. I know that there is a laptop computer in front of me. I know that I am typing away on that computer with those very hands, *my* hands. I know that I am in funky coffeehouse. I know that the music of Nickel Creek is playing in the background. And, sadly, I know that my coffee cup is almost empty. These are just some of my many epistemic accomplishments. You know a lot, too. You know that you have eyes. You know that you are seeing a journal article (or a computer screen image of a journal article) with those eyes. You know that the article you are reading is written in English. Together, we know a lot. At least, we think we do, until we encounter the skeptic.

The skeptic contends that all of the above knowledge ascriptions are false. To convince us of our ignorance, the skeptic puts forth some skeptical hypothesis *H* such that if *H* were true, then (i) our evidential situation would be phenomenologically indistinguishable from our current evidential situation, and yet, (ii) all our commonsense perceptual beliefs would be false. The skeptic then argues that since we don't know that *H* is false, we don't know that the beliefs of commonsense are true. We might be being systematically deceived by the malevolent machinations of Descartes's demon or by the brain-distorting effects of Lehrer's Googols.<sup>1</sup> Or, consider the ever popular brain-in-a-vat [BIV] hypothesis, according to which I am a disembodied brain floating in a vat of nutrient hooked up to a sophisticated computer that is producing coherent experiences in me by stimulating my sensory cortex in exactly the same way it would have been stimulated were I perceiving normally. Since my being a *handless* BIV entails that I don't have hands, the skeptic argues as follows:

**BIV arguments****BIV 1**

1. I don't know that I'm not a handless BIV.
2. If I don't know that I'm not a handless BIV, then I don't know that I have hands.  
Therefore,
3. I don't know that I have hands.

Not wanting you to miss out on all the fun, the skeptic also argues as follows:

## DIV 2

1. You don't know that you're not an eyeless BIV.
2. If you don't know that you're not an eyeless BIV, then you don't know that you have eyes.  
Therefore,
3. You don't know that you have eyes.

These arguments are valid and their premises seem unassailable, and yet their conclusions strike us as absurd. Of course, accepting their premises while rejecting their conclusions commits us to what Stewart Cohen calls the inconsistent triad: 1, 2, and  $\sim 3$ . Our options aren't appealing. Rejecting these arguments' second premises commits us to two rather unsavory abominable conjunctions: "I know that I have hands, but I don't know that I'm not a *handless* BIV" and "You know that you have eyes, but you don't know that you're not an *eyeless* BIV"; and rejecting these arguments' first premises requires our knowing that we are not BIVs, but how could we know that? And yet, surely, I know that I have hands (with which I am currently typing), and you know that you have eyes (with which you are currently reading!). We're in the skeptical fly bottle, alright. Something has to go, but what?

That's where contextualism comes in. In true Wittgensteinian fashion, contextualists promise to show us the way out of the fly bottle. But, they stress, not just any way out of the bottle will do: One can't just reject a premise and release the flies; a *satisfactory* solution to the skeptical problem must explain why we feel ensnared in the first place. As Keith DeRose puts it:

In seeking a solution to this puzzle, we should seek an explanation of how we fall into this trap in the first place, and not settle for making a simple choice among three distasteful ways out of the trap. We must explain how two premises that together yield a conclusion we find so incredible can themselves seem so plausible to us (DeRose 1995, p. 3).<sup>2</sup>

In particular, contextualists insist that an adequate solution to the skeptical problem must meet the following three *desiderata*: (i) It must retain epistemic closure,<sup>3</sup> (ii) it must explain the intuitive force of skeptical arguments by explaining why their premises initially seem so compelling, and (iii) it must account for the truth of our commonsense judgment that we do in fact possess lots of ordinary knowledge. The key to such a solution, according to contextualism, is recognizing that the semantic standards for "knows" vary from context to context such that in skeptical contexts the skeptic's

premises are true and so is her conclusion; but in ordinary contexts, her conclusion is false and so is her first premise. Of course, there are embellishments. In what follows, I will provide those embellishments, for the devil is in the details. I will first explain the main tenets — the defining core — of the position that has come to be known as *epistemic contextualism*. Epistemic contextualism is viewed by some as a kind of panacea for nearly all of our epistemic ills. At the very least, it has been invoked to cure several persistent epistemological hangovers, including the lottery paradox and the Gettier problem, but its single biggest selling point by far is the apparent ease with which it can resolve the skeptical paradox. If it fails to deliver in this latter regard or if there is a preferable noncontextualist solution, then one of the primary motivations for contextualism collapses. In what follows, I present the contextualist 'solution' to the skeptical problem. I then point out a number of counterintuitive costs associated with contextualism. If we can provide an alternative noncontextualist solution to the skeptical problem that meets the contextualist's adequacy constraint on such solutions, while avoiding contextualism's counterintuitive results, then the primary motivation for contextualism will have been undermined. In the paper's final sections, I propose such a noncontextualist solution to the skeptical problem. Since epistemic contextualism is rooted in semantic contextualism, let's begin our discussion there.

## 2. SEMANTIC CONTEXTUALISM

In "Scorekeeping in a Language Game", David Lewis rightly observes that when interpreting people's utterances, we employ various context-sensitive *rules of accommodation* that facilitate effective communication (Lewis 1979). Lewis also stresses that contexts can shift very quickly, and we routinely succeed in accommodating these shifts. Here's an embellished version of his famous cat case. We are sitting in a garden in Mainz, Germany drinking a beer, and you've just set your beer on the park bench next to you. I'm telling you about one of my cats back in Illinois, and I say: "The cat loves to chase string. When she judges that it's been too long since we've played string, the cat picks up the string in her mouth and brings it over to me with a 'Play-with-me-dammit!' look in her eye. Look out, the cat is about to spill your beer!" You immediately interpret my latter use of "the cat" to refer to the stray cat that has just jumped up on the bench next to you. Your beer safely in hand, I continue: "I'm sure when I return home,

the cat will greet me at the door, string in mouth," You accommodate each of my uses of "the cat". You realize that my first two uses refer to my cat in Illinois, my third use refers to the stray cat on the bench, and my fourth use once again refers to my cat back home. So far so good. Semantic contextualism is a truism about language.

In the case just described, speaker intentions [my intentions] determined which cat I was referring to when I used the expression "the cat". Context did not determine which cat I was referring to, I did. To be sure, context played a role in allowing you to ascertain which cat I was referring to. Since a cat in Illinois is unlikely to spill a beer in Mainz, you rightly interpreted me to be talking about the stray cat in Mainz, and since a cat in Mainz won't likely greet me when I return to Illinois, you rightly interpreted me to be talking about my own cat again. Epistemic contextualism takes as its starting point the truism that language is highly context-sensitive in just the way that semantic contextualists maintain.

### 3. EPISTEMIC CONTEXTUALISM

While epistemic contextualists differ on the details of their theories, e.g. some work within the relevant alternatives framework (see Cohen 1988) and others embrace the subjunctive conditionals approach to knowledge (see DeRose 1995), there are certain core defining features on which all epistemic contextualists agree.

#### 3.1. *The Metalinguistic Turn*

First, contextualists maintain that there is no correct context-independent standard of knowledge. Consequently, there is no context-independent fact of the matter as to whether or not *S* knows that *p*. Since there is no fact of the matter as to whether or not *S* knows that *p* outside a context of ascription, they maintain that we epistemologists should drop all talk about whether or not *S* knows that *p*. Our focus, instead, should be on whether sentences of the form "*S* knows that *p*" are true in some specified context of ascription.

#### 3.2. *Ascriber-Sensitive Contextualism*

The second defining feature of epistemic contextualism is that it is the context of the *knowledge-ascriber*, not the context of the

knowledge-ascriber, that determines which standards of knowledge are operant in her assertion. The truth conditions for ascriber *A*'s assertion "*S* knows that *p*" are determined by *A*'s context, not by *S*'s context. Consequently, for a given cognitive subject *S*, proposition *p*, and time *t*, it is possible for an ascriber *A1* to truthfully assert "*S* knows that *p*" at *t*, while another ascriber *A2* truthfully asserts "*S* does not know that *p*" at *t*, provided, e.g., that *A1* is in a low standards context and *A2* is in a high standards context. Given the different contexts of utterance, the proposition expressed by *A2*'s utterance is *not* the negation of the proposition expressed by *A1*'s utterance. So, *A1* and *A2* need not be disagreeing.

While the standard form of epistemic contextualism is *ascriber-sensitive contextualism*,<sup>4</sup> as outlined above, it must be stressed that - unlike Lewis's cat example above - it's not up to the speaker/ascriber to decide what sense of "knows" she is using. If the ascriber finds herself in a skeptical context - perhaps she's just been discussing deceptive demons in an epistemology class and is still worrying about them - the word "knows" *in her mouth* will refer to high standards knowledge whether she intends it to or not. Even if she wants to use "knows" in its low standards sense, she won't be able to as long as she is contemplating demons. It is the *context* of the ascriber, not the ascriber herself, that determines what standards of knowledge are operant in her assertion. As we shall subsequently see, the context of an ascriber *A* can be impacted and altered by the content of proposition *p* in *A*'s assertion "*S* knows that *p*".

### 3.3. *Asymmetrical Standards Adjustments*

Epistemic contextualism differs from semantic contextualism in another important respect, as well. With Lewis's cat example, it was relatively easy to shift from Illinois cat to Mainz cat, and it was equally easy to shift from Mainz cat back to Illinois cat. But in the case of epistemic contextualism, there is an asymmetry in the amount of ease with which one can move from one use of "knows" to another. Epistemic contextualists maintain that it is easy to move from a low standards context of "knows" to a high standards context of "knows", but they insist that it is very difficult to move from a high standards context back down to a low standards context. Once the epistemic standards have been raised, they tend to stay raised.<sup>5</sup>

## 4. CONTEXTUALISM TO THE RESCUE: SOLVING THE SKEPTICAL PROBLEM

We are now in a position to see how epistemic contextualism purports to solve the skeptical problem. As we go about our day moving from one mundane task to the next, I know that I have hands and you know that you have eyes. Better to state it metalinguistically: In an ordinary context, my knowledge self-ascription "I know that I have hands" is true, and so is my knowledge attribution "You know that you have eyes", because in an ordinary context, the operant standards of knowledge tend to be low. [In such an ordinary context, your knowledge self-ascription asserting that you know you have eyes (i.e. "I know that I have eyes" *in your mouth*) would be true for similar reasons.]

Then, we meet a skeptic who presents us with the aforementioned BIV arguments. Prior to our encounter with the skeptic, the BIV hypothesis was not a relevant alternative, and hence, it was perfectly proper of us to ignore it. However, by the very act of mentioning the BIV hypothesis, the skeptic makes the BIV alternative relevant and improperly ignored, thereby in effect raising the standards of knowledge on us. Because I can't rule out the BIV alternative, the statement "I don't know I'm not a BIV" is true *in my mouth* in that context, and as long as the BIV hypothesis remains relevant for me, "I don't know that I have hands" is also true for me in that context. Likewise for you. Because you can't rule out the BIV hypothesis either, the sentence "I don't know that I'm a BIV" is true *in your mouth* as well, and as long as it remains true for you, the sentence "I don't know I have eyes" is also true *in your mouth*. That, according to contextualists, is why we feel so threatened by skeptical arguments. By presenting us with such arguments, the skeptic manipulates the semantic standards of "knows" and moves us from an ordinary context in which low standards of knowledge are in effect to a skeptical context in which high standards are in effect; and we realize that, given the newly operant high standards of knowledge, we don't know the things we ordinarily take ourselves to know.

But, contextualists insist, the skeptic doesn't win the day, because like Hume once we leave our skeptical closet, return to our day-to-day affairs, and forget about the skeptic's challenge, our ordinary low standards become operant again, and our knowledge claims in that context are true. In fact, the skeptic was never able to show that we lacked low standard ordinary knowledge, because by presenting the skeptical argument the skeptic changed the topic from ordinary knowledge, where demons and BIV-scenarios are properly ignored, to

high standards knowledge, where these alternatives can't be ignored. And so, contextualists maintain, we have a solution to the skeptical puzzle that meets the contextualist's three stated *desiderata*. By the very act of asserting the BIV argument, the skeptic makes salient a skeptical alternative that we cannot rule out. In that context, the skeptic's premises are true, and so is her conclusion. That explains why we find the skeptic's argument so troubling. But when we leave our skeptical closets and forget about the skeptic's alternatives, these alternatives are no longer salient, and our ordinary low standards once again become operant, thus explaining why we remain convinced that we know that we have hands and eyes. Closure holds in any given context.

#### 5. THE PRICE OF SUCCESS/THE COST OF CONTEXTUALISM

Despite the apparent ease with which it handles the skeptical problem, epistemic contextualism has numerous counterintuitive consequences. Here are a few of its unsavory results:

##### 5.1. *Hume*

Contextualism entails that *from Hume's own perspective*, relative to the standards operant for knowledge self-ascriptions, Hume "knew" more in the bar when he was three sheets to the wind, than he did when he was carefully reflecting in his closet. Lewis makes the point as follows: "... when we do epistemology, and we attend to the proper ignoring of possibilities, we make knowledge vanish. First we do know, then we do not" (Lewis 1996, p. 566). But that's not quite right. Strictly speaking, contextualism does not entail that any low standard knowledge is lost. Hume does not lose his low standard knowledge when he is entertaining skeptical reflections in his study, because *some other person* in a low standard context could still truthfully ascribe low standards knowledge to Hume, when Hume is in his closet.<sup>6</sup> But I am not denying that there is *some perspective* from which one can truthfully attribute knowledge to Hume. The present worry is that *Hume* couldn't ascribe such knowledge to himself when he is in his skeptical closet. Contextualism implies that, *from Hume's own perspective*, he can't truthfully ascribe knowledge to himself when he is entertaining skeptical thoughts in his closet, but he can truthfully ascribe knowledge to himself when he is cavorting in the bar. Why? Because when Hume was contemplating various skeptical



hypotheses in his study, these alternatives were salient for him and could not properly be ignored. Of course, when he was making merry in the bar and ladies, not demons, were on his mind, evil demon scenarios were no longer relevant, and so, his slurred knowledge self-ascriptions were true. This is the sense in which contextualism entails that, *from Hume's own perspective*, he knew more in the bar when he was semi-inebriated than he did when he wrote *A Treatise of Human Nature* while carefully reflecting in his study. Contextualism also entails that when Hume was in the high standards context of his skeptical closet, he could not know that he still had low standards knowledge.<sup>7</sup> One wonders how a person could write such an important and influential philosophical treatise, while, *from his own perspective*, knowing absolutely nothing (i.e. while, *from his own perspective*, all of his knowledge self-ascriptions are false). It strikes me as implausible in the extreme to think that more of Hume's knowledge self-ascriptions were true when he was half drunk in the bar than were true when he was soberly reflecting as carefully as possible on his epistemic situation. It also strikes me as implausible that Hume could have written a work of such lasting philosophical importance, while *from his own perspective* knowing absolutely nothing at all. How did he even manage to find a quill pen in his closet of absolute ignorance?

### 5.2. *Therapy 1*

I often find that when I go to philosophy conferences, esp. epistemology conferences, most of the participants know a lot more than I do. Not wanting to feel inferior, whenever I encounter these daunting intellectuals, I just contemplate Descartes's evil demon, and just like that, I can *truthfully* assert: "These people don't know more than me. In fact, they don't know anything at all." I feel better. Contextualist therapy at work.

### 5.3. *Therapy 2*

Sometimes I meet a particularly smug philosopher in need of therapy himself. He has not so subtly made it clear that he takes himself to know far more about the subject at hand than I do. In light of my recent therapy session above, I can truly assert "He doesn't know anything at all", and I want to enlighten him to this truth. So, I walk up to him and say: "You might be a BIV or a victim of wholesale

demonic deception. Deal with your total ignorance". Just like that, I render all of his knowledge self-ascriptions false. If he claims to know any proposition in his newly acquired demon-context, his claim will be false. Now we're epistemic equals, and once again, I feel good.

But surely something is amiss. One can't make the sentence "No one knows anything" true as easily as Therapy 1 suggests. Nor can one convince people of the truth of sentences ascribing a total state of ignorance to them, as simply Therapy 2 suggests.

#### 5.4. *Unspeakable knowledge*

The contextualist is also committed to the view that there are many true propositions that we *know* are true, but that are such that it is impossible for us to truthfully assert that we know them. Consider the true proposition *that I am not a BIV*. According to our contextualist-sanctioned ordinary low standard knowledge, I know that this proposition is true, but I cannot truthfully assert "I know that I am not a BIV" , because in the very act of making such an assertion, I inevitably raise the standards for knowledge to a level so high that I no longer satisfy those standards with respect to the proposition *that I am not a BIV*. Such knowledge, which we do in fact possess, is not only unspeakable, it is unthinkable! As soon as we think about our not being BIVs, we cease to know we're not BIVs.<sup>8</sup>

#### 5.5. *It's Hard to be a Consistent Contextualist*

If there ever were a skeptical context, this is it. I've been discussing deceptive demons, brain-distorting Googols, and BIV-scenarios, and my doing so has forced you to contemplate these skeptical alternatives. Before I mentioned these alternatives, they were irrelevant and properly ignored. But now that we're thinking about them, they are relevant and improperly ignored. So, if contextualism is correct, right now *in the present context*, my ascription "You don't know that you have eyes" is true, and if you protest "But I do know that I have eyes", your assertion is false. Since we are in a skeptical context, *in your mouth*, the utterance "I know that I have eyes" is false. Moreover, not only is the utterance false, so is the thought. If right now you are thinking, "How silly! Of course, I know I have eyes!" , your belief is false, according to contextualism. By mentioning demons, Googols, and hallucination-producing vats, I've raised the standards for knowing, and you can't just lower them again at will. It takes awhile to forget what has been said, and until you do, *from your own perspective*,

you won't know that you have eyes [i.e. the knowledge ascription "I know that I have eyes" *in your mouth* will be false]. But, I daresay, there is not a single person *reading* this article who doubts that she knows she has eyes. Not even the contextualists among us would think that the sentence "I know I have eyes" is false *in their mouths/minds* while they are reading this article, despite my having just mentioned the BIV hypothesis. Despite the fact that contextualism entails that, in the present demon-infested context, no one reading this article knows that s/he has eyes, everyone reading it, including contextualists, continues to believe that s/he knows that s/he has eyes.

### 5.6. *Vanishing Force*

There is one phenomenon in particular that contextualism *in principle* cannot explain, namely, the widely recognized phenomenon that skeptical arguments tend to lose their force as we contemplate them time and again. Contextualism predicts just the opposite. It predicts that every time we contemplate a BIV argument, a demon argument, or any other skeptical argument, we will once again find ourselves in the throws of skepticism. It predicts that whenever we are presented with such an argument, we will conclude "I really don't know what I thought I knew", but we conclude no such thing. The contextualist can't explain why the first time people encounter an argument like the BIV argument, they are often filled with epistemic angst, but as they study more epistemology, they cease to be moved by such arguments. To its credit, my noncontextualist solution does account for the vanishing force of skeptical arguments.

## 6. TOWARDS A NONCONTEXTUALIST RESOLUTION OF THE SKEPTICAL PROBLEM

I accept the contextualist's challenge that an adequate solution to the skeptical problem must do three things: (i) It must retain epistemic closure, (ii) it must explain the intuitive force of skeptical arguments by explaining why their premises initially seems so compelling, and (iii) it must account for the truth of our commonsense judgment that we do in fact possess lots of ordinary knowledge. Since I'm giving a noncontextualist solution to the skeptical problem that retains epistemic closure while maintaining that we do possess a great deal of knowledge, I must attack premise 1, but in the course of doing so, I must also explain why premise 1 *initially* strikes us as being so

plausible. Often when presented with a paradoxical argument like the BIV argument, the action is really taking place at the level of unstated assumptions. Why is it that we are initially so convinced by the skeptic's first premise? How does the skeptic persuade us to accept that premise? As already noted, she does so by presenting us with a detailed skeptical hypothesis *H* such that if *H* were true, then (i) our evidential situation would be phenomenologically indistinguishable from our current evidential situation, and yet, (ii) all our common-sense perceptual beliefs would be false. She then argues that the mere *possibility* of *H* is sufficient to prevent us from having any perceptual knowledge at all. It is important that *H* be possible. No one was ever moved to the brink of skepticism by the skeptical hypothesis: "You might be being deceived by a malevolent roundsquare". In presenting her skeptical hypothesis, the skeptic asserts that it is *possible* that I am a BIV. Starting with this premise, she then argues in one of two ways. Either she uses the *Argument from Possibility 1* to defend the first premise of BIV1 as follows:

**Argument from Possibility 1 (AP1)**

1. It is possible that I am a handless BIV.
  2. If it's possible that I am a handless BIV, then I don't know that I'm not a handless BIV.
- Therefore,
3. I don't know that I'm not a handless BIV.

Or, she uses the *Argument from Possibility 2* to argue straight away for the conclusion that I don't know that I have hands:

**Argument from Possibility 2 (AP2)**

1. It is possible that I am a handless BIV.
  2. If it's possible that I am a handless BIV, then it's possible that I don't have hands.
  3. If it's possible that I don't have hands, then I don't know that I do have hands.
- Therefore,
4. I don't know that I have hands.<sup>9</sup>

Like its BIV1 counterpart, AP1 looks valid and its premises certainly seem to be true. So, we seem to have a good argument for AP1's conclusion, an argument which makes premise 1 of BIV1 *prima facie* plausible. AP2 also appears valid by repeated instances of *modus ponens*, and its premises likewise seem true. So, AP2 seems

to provide compelling reason to think that I don't know that I have hands. In what follows, I will argue that AP2 either begs the question or equivocates, and either way, fails to give us a good reason to accept AP2's skeptical conclusion. Since similar arguments *mutatis mutandis* will show that AP1 also either begs the question or equivocates, and so, fails to provide a good reason for accepting premise 1 of BIV1, I won't rehearse those arguments. To set the stage for my argument, I must briefly discuss epistemic possibility.

### 6.1. *Epistemic Possibility: Two Cases*

Consider the following case:

#### **Philosopher Bob**

Ordinary Joe and Philosopher Bob are sitting along the lakeshore in Chicago discussing Jim Java's whereabouts.

*Joe:* I haven't seen Jim Java in a few days. Perhaps, he is in New Orleans. He told me he was going to take a trip there soon, just so he could drink chicory coffee in the French Quarter.

*Bob:* It's not possible that Jim Java is in New Orleans. I just saw him at the Bump and Grind Coffeehouse twenty minutes ago.

Bob intends to assert something true. But on a metaphysical reading, his modal claim is false,<sup>10</sup> because there is a  $\mu$ -possible [metaphysically possible] world where Jim was at the Bump and Grind twenty minutes ago, but also where intercity transportation is remarkably efficient, so efficient in fact that one can get from Chicago to New Orleans in 10 min. Of course, Bob knows that that world is not actual. In the actual world, it takes twenty minutes just to get a cab in Chicago, another thirty minutes to get to O'Hare Airport, and an hour to get through security. Given what Bob knows, Jim simply couldn't be in New Orleans. He couldn't have even made it to O'Hare in so little time. Bob's background knowledge entails that Jim is not in New Orleans, and Bob recognizes that entailment. Given Bob's knowledge, it is not e-possible [epistemically possible] for Jim to be in New Orleans. That is why Bob asserts what he does. Bob is making a true epistemic modal claim, not a false metaphysical one.

Consider a second case due to Kripke:

**Goldbach's Conjecture**

Goldbach's Conjecture (GC) is the mathematical conjecture that every even number greater than 2 is the sum of two primes. While no counter-instance to GC has ever been found, no one has ever demonstrated that GC is true. Since GC is a mathematical proposition, it has its truth-value necessarily. So, *if* GC is true (as is widely believed), then it is *necessarily* true; and yet, since it has never been proven, it seems true to say that it *might* be false. Such is the nature of conjectures. They might be false. How *might* a necessarily true proposition be false? In what sense, is it *possible* for a necessarily true proposition to be false? Kripke's answer is, "in the epistemic sense." He rightly observes that the "might" and the "possible" are being used in an epistemic sense merely to express our present ignorance of the truth value of GC (Kripke 1980, pp. 36–38).

6.2. *A Stipulative Account of E-Possibility*

As Ian Hacking has observed, certain "occurrences of *possible* can be modified by many adverbs of the form  $\phi$ -ly: technically, economically, theoretically, medically, metaphysically, humanly" (Hacking 1975, p. 325). We can understand  $\phi$ -possibility ascriptions using the following  $\phi$ -possibility schema:

$\Diamond_{\phi} p$   $p$  is  $\phi$ -ly possible for  $S$  iff nothing of a  $\phi$ -al sort  $\phi$ -ly precludes the truth of  $p$ .

According to this schema,  $p$  is logically possible for  $S$  iff nothing of a logical sort *logically precludes* the truth of  $p$ ; and  $p$  is physically possible iff nothing of a physical sort *physically precludes* that  $p$ . Similarly,  $p$  is epistemically possible for  $S$  iff nothing of an epistemic sort *epistemically precludes* the truth of  $p$ . Let us, therefore, turn our attention to the notion of "epistemic preclusion" since it holds the key to understanding epistemic possibility.

6.3. *Epistemic Preclusion*

There are two ways that something known by  $S$  can *epistemically preclude* the truth of  $p$  for  $S$ . The first and most obvious way for  $p$  to be epistemically precluded for  $S$  is for  $S$  to know that  $\sim p$ . Let's call this "direct e-preclusion" and define it as follows:

$D_1$   $p$  is directly e-precluded for  $S$  at  $t$  iff  $S$  knows that  $\sim p$  at  $t$ .

To see that there is another way for something  $S$  knows to epistemically preclude  $p$  for  $S$  consider the following scenario.

**Simone**

I am lying awake in bed. Simone, one of my cats, is curled up against my side purring loudly. I know that Simone is in the bedroom, but I do not know that she is not in the kitchen, because I haven't hothered to form the belief that she is not in the kitchen. Still, my knowledge that she is in the bedroom, together with my background knowledge, self-evidently entails that she is not in the kitchen, i.e. this entailment is one I would immediately recognize, were I to consider it. My knowledge that Simone is in the bedroom indirectly e-precludes her being in the kitchen. Call this *indirect e-preclusion*:

$D_2$   $p$  is indirectly e-precluded for  $S$  at  $t$  iff (i)  $S$  does not know that  $\sim p$  at  $t$ , but (ii)  $S$  could come to know that  $\sim p$  at  $t$ , strictly on the basis of propositions  $S$  knows at  $t$ .<sup>11</sup>

With these definitions in hand, we can define *e-possibility* as follows:

$D_3$   $p$  is e-possible for  $S$  iff  $p$  is neither directly nor indirectly e-precluded for  $S$ .

Or equivalently:

$D'_3$   $p$  is e-possible for  $S$  at  $t$  iff (i)  $S$  does not know that  $\sim p$  at  $t$ , and (ii)  $S$  could not come to know that  $\sim p$  at  $t$ , strictly on the basis of propositions  $S$  knows at  $t$ .<sup>12,13</sup>

$D'_3$  yields the right results where Philosopher Bob is concerned, for it entails that it's *not* e-possible for Bob that *Joe is in New Orleans*, because Bob knows that (i) Joe was at the Bump and Grind Coffeehouse in Chicago twenty minutes earlier and he also know that (ii) in the actual world it takes more than twenty minutes to get from Chicago to New Orleans; and Bob recognizes that (i) and (ii) entail that Joe is not in New Orleans.<sup>14</sup>  $D'_3$  also yields the right results where Goldbach's conjecture is concerned. I do not know that GC is true, and nothing I currently know self-evidently entails that GC is true. [Neither I nor any expert mathematician has been able to see how to derive GC.] So,  $D'_3$  entails that  $\sim$ GC is e-possible for me and *ipso facto* that it is e-possible for me that GC is false.

6.4. *The "Possibility" of the BIV Hypothesis*

One problem with the *Argument from Possibility* (AP) should be obvious—it does not specify the kind of possibility it is employing in its premises. To help us sort through the various readings of AP2, let's symbolize it as follows:

- AP2     1.  $\Diamond b$   
           2.  $\Diamond b \rightarrow \Diamond \sim h$   
           3.  $\Diamond \sim h \rightarrow \sim Kh$   
 $\therefore$      4.  $\sim Kh$ <sup>15</sup>

Perhaps, the skeptic hopes to establish the skeptical conclusion by appealing to the mere metaphysical possibility of my being a BIV as follows:

- MAP2    1.  $\Diamond_{\mu} b$   
           2.  $\Diamond_{\mu} b \rightarrow \Diamond_{\mu} \sim h$   
           3.  $\Diamond_{\mu} \sim h \rightarrow \sim Kh$   
 $\therefore$      4.  $\sim Kh$ <sup>16</sup>

MAP2 is valid, but it is not sound. Premise 3 is false. We cannot, generally speaking, derive epistemic conclusions from purely metaphysical premises.<sup>17</sup> And we certainly cannot derive *S*'s ignorance of *p* from the mere  $\mu$ -possibility of  $\sim p$ , for, as we have already seen with Philosopher Bob, the  $\mu$ -possibility of  $\sim p$  does not imply that *S* does not know that *p*. The  $\mu$ -possibility that Jim Java is in New Orleans does not imply that Philosopher Bob does not know that Jim is not in New Orleans. Bob does know that Jim is not in New Orleans, even though, being a philosopher, he also knows that it is  $\mu$ -possible that Jim is there.<sup>18</sup>

If the skeptic wishes to derive the epistemic conclusion that I don't know I'm not a BIV from AP2, then she must appeal to epistemic possibility throughout her argument:

- EAP2    1.  $\Diamond_e b$   
           2.  $\Diamond_e b \rightarrow \Diamond_e \sim h$   
           3.  $\Diamond_e \sim h \rightarrow \sim Kh$   
 $\therefore$      4.  $\sim Kh$ <sup>19</sup>



### 6.5. *Exploring the Epistemic Possibility of the BIV Hypothesis*

Unlike MAP2, EAP2 presents us with a genuine skeptical problem – a paradox of its own, for it looks valid and its premises seem to be true, and yet, its conclusion strikes most epistemologists as false.<sup>20</sup> That premise 3 is true can be demonstrated as follows:  $D_3^1$  entails the following necessary condition for e-possibility: *p* is e-possible for *S* only if *S* does not know that  $\sim p$ .<sup>21</sup> Formally, the condition can be stated as follows:  $(p)(0_{ep} \rightarrow \sim K \sim p)$ .<sup>22</sup> Since premise 3 is an instance of this general truth, premise 3 is clearly true. Premise 2 is also true. The e-possibility of my being a handless BIV does imply the e-possibility of my having no hands; for if nothing I know e-precludes my being a handless BIV, then nothing I know e-precludes my having no hands. Given EAP2's validity and the truth of its premises 2 and 3, it follows that if premise 1 of EAP2 is true, then EAP2's conclusion 4 must also be true. And premise 1 looks true – it certainly seems e-possible that I am a handless BIV. Of course, it also seems clear that I know that I have hands. After all, I have a reliably produced, perceptually justified true belief that I have hands, and there is no Gettier funny-business going on. The paradox generated by EAP2 then is this: The BIV hypothesis is e-possible, and yet we know things incompatible with its e-possibility. How can that be? In what follows, I will resolve the paradox in a way that allows us to retain our philosophical intuition that the BIV hypothesis is e-possible, while also allowing us to retain our commonsense intuition that we do have knowledge of the external world around us. I will argue that EAP2 fails to provide us with a reason to think that external world skepticism is true, because it either rests on an equivocation or it begs the question.

### 6.6. *Two Kinds of E-Possibility*

I submit that, even in EAP2, there lurks a hidden ambiguity, because there is more than one kind of e-possibility. Let me explain. One of the reasons it is difficult to get an intuitive handle on e-possibility is because our intuitive e-possibility assessments are split along the same infallibilistic/fallibilistic lines as our ordinary epistemic evaluations.<sup>23</sup> We typically relativize our e-possibility assignments to the propositions we *fallibly* know [ $know_f$ ], but occasionally, we make our e-possibility assignments relative to the propositions we *infallibly* know [ $know_i$ ].<sup>24</sup> To avoid conflating our fallibilistic e-possibility

assessments with our infallibilistic e-possibility assessments,  $D'_3$  must be revised as follows:

- $D'_{3f}$   $p$  is e-possible<sub>f</sub> for  $S$  at  $t$  iff (i)  $S$  does not know<sub>f</sub> that  $\sim p$  at  $t$ , and (ii)  $S$  could not come to know<sub>f</sub> that  $\sim p$  at  $t$ , strictly on the basis of propositions  $S$  knows<sub>f</sub> at  $t$ .
- $D'_{3i}$   $p$  is e-possible<sub>i</sub> for  $S$  at  $t$  iff (i)  $S$  does not know<sub>i</sub> that  $\sim p$  at  $t$ , and (ii)  $S$  could not come to know<sub>i</sub> that  $\sim p$  at  $t$ , strictly on the basis of propositions  $S$  knows<sub>i</sub> at  $t$ .<sup>25</sup>

Once we recognize the distinction between fallibilistic and infallibilistic e-possibility, EAP2 itself turns out to be multiply ambiguous between a purely infallibilistic reading [EAP2i], two mixed readings [EAP2<sub>x1</sub> and EAP2<sub>x2</sub>], and a purely fallibilistic reading [EAP2f]:

EAP2i	1. $\Diamond_{ei}b$	EAP2 <sub>x1</sub>	1. $\Diamond_{ei}b$
	2. $\Diamond_{ei}b \rightarrow \Diamond_{ei}\sim h$		2. $\Diamond_{ei}b \rightarrow \Diamond_{ei}\sim h$
	3. $\Diamond_{ei}\sim h \rightarrow \sim K_i h$		3. $\Diamond_{ei}\sim h \rightarrow \sim K_i h$
$\therefore$	4. $\sim K_i h$ <sup>26</sup>	$\therefore$	4. $\sim K_i h$ <sup>27</sup>
EAP2 <sub>x2</sub>	1. $\Diamond_{ei}b$	EAP2f	1. $\Diamond_{ef}b$
	2. $\Diamond_{ei}b \rightarrow \Diamond_{ef}\sim h$		2. $\Diamond_{ef}b \rightarrow \Diamond_{ef}\sim h$
	3. $\Diamond_{ef}\sim h \rightarrow \sim K_i h$		3. $\Diamond_{ef}\sim h \rightarrow \sim K_i h$
$\therefore$	4. $\sim K_i h$	$\therefore$	4. $\sim K_i h$

Neither mixed reading is plausible. EAP2<sub>x1</sub> identifies an equivocal reading of EAP2 and is simply invalid. As for EAP2<sub>x2</sub>, its second premise is false. The fact that none of the propositions I know<sub>i</sub> e-precludes that I am a BIV does not entail that none of the propositions I know<sub>f</sub> e-precludes my not having hands. Why? Because I might know<sub>f</sub> that *I have hands* and that knowledge<sub>f</sub> would fallibilistically directly e-preclude that I don't have hands without knowing<sub>i</sub> that I am not a BIV and without knowing<sub>i</sub> anything that entails or justifies<sub>i</sub> me in believing that I'm not a BIV.

As for EAP2i, while it is clearly sound, it isn't of much philosophical interest. First to its soundness. EAP2i is valid, and its second and third premises are true for reasons analogous *mutatis mutandis* to those offered when discussing premises 2 and 3 of EAP2 above. That leaves premise 1 of EAP2i to consider. It is generally acknowledged that we have very little in the way of knowledge<sub>i</sub> for our evidence rarely entails that for which it is evidence. We may know<sub>i</sub> a few *cogito*

propositions, but that's about it. Given the little, if anything, that we know<sub>i</sub>, very few propositions, if any, are infallibilistically e-precluded for us. Suppose, for example, that I possess no knowledge<sub>i</sub> whatsoever. Then, no propositions are infallibilistically e-precluded for me, and so every proposition is e-possible<sub>i</sub> for me. If, on the other hand, I do possess *cogito* knowledge<sub>i</sub> of my own existence, then that knowledge<sub>i</sub> infallibilistically e-precludes my own nonexistence for me. But my *cogito* knowledge<sub>i</sub> does not infallibilistically e-preclude my being a BIV, because I cannot justifiably<sub>i</sub> infer my nonenvattedness from the few *cogito* propositions I know<sub>i</sub>. Since my being a BIV is not infallibilistically e-precluded for me, my being a BIV is e-possible<sub>i</sub> for me, just as premise 1 of EAP2i asserts.<sup>28</sup> Hence, EAP2i is sound. I don't know<sub>i</sub> that I have hands.

The reason EAP2i is uninteresting is because I don't need to contemplate the e-possibility<sub>i</sub> of far-fetched BIV hypotheses to realize that I lack knowledge<sub>i</sub> that I have hands. Presumably, my current visual and tactile experiences are what justify me in believing that I have hands, and it is obvious that those experiences *do not entail* that I have hands, for I can have those same experiences as a result of dreams, mushroom-induced hallucinations, virtual reality machines, phantom limb experiences, etc. Therefore, I am not justified<sub>i</sub> in believing that I have hands. Since I lack justification<sub>i</sub> for believing that I have hands, and since justification<sub>i</sub> is necessary for knowledge<sub>i</sub>, it follows that I lack knowledge<sub>i</sub> that I have hands. So, we do not need to appeal to EAP2i to establish such a conclusion.

The only interesting version of EAP2 is EAP2f, since it is the only version which threatens to undermine our ordinary fallibilistic knowledge of the objects around us. Like EAP2i, EAP2f is valid and its second and third premises are true, again for reasons analogous *mutatis mutandis* to those offered in support of EAP2's premises 2 and 3. The problem with EAP2f is that the skeptic is in no position to assert that its first premise is true, for *suppose* that I know<sub>f</sub> that I have hands (h). Then I know<sub>f</sub> a proposition – namely, h – that entails that I am not a handless BIV.<sup>29</sup> And since I have the cognitive power needed to grasp this entailment, my knowing<sub>f</sub> that h would fallibilistically (indirectly) e-preclude for me my being a handless BIV and would thus render the BIV hypothesis e-impossible<sub>f</sub> for me (regardless whether I have actually noticed the entailment or not). So, the skeptic can only rationally assert EAP2f's first premise – that it is e-possible<sub>f</sub> for me that I am a handless BIV – if she *assumes* the truth of EAP2f's conclusion. Since the skeptic cannot rationally assert EAP2f's first premise without assuming that I lack knowledge<sub>f</sub>

that I have hands, she cannot assert premise 1 without assuming the very thing in question. Granted, if the skeptic could give an independent reason for thinking EAP2f's first premise true – a reason that did not make reference to the truth of EAP2f's conclusion, then EAP2f would not beg the question; but she can't because  $e$ -possibility<sub>f</sub> is analyzed in terms of knowledge<sub>f</sub>. Consequently, EAP2f effectively begs the question, because to be rationally entitled to assert premise 1 of EAP2f, the skeptic must *first* be rationally entitled to assert that  $\sim K_f h$ . Perhaps the skeptic can provide some other argument for  $\sim K_f h$ , which she can then use to establish  $\sim K_f h$  and *ipso facto* EAP2f's first premise. But then, it is this other argument – not EAP2f – that is doing all the skeptical work. Any argument A1 such that one must first establish *the conclusion* of A1 via some second argument A2 *before* one can rationally assert the premises of A1 is itself worthless in establishing the conclusion of A1. EAP2f is such an argument. In order for the skeptic to rationally assert premise 1 of EAP2f, she must first prove the truth of EAP2f's conclusion with a different argument, thereby rendering EAP2f superfluous.<sup>30</sup>

#### 6.7. Undermining AP2 and Resolving the EAP2 Paradox

An adequate solution to the skeptical problem must not only explain where the skeptic's argument goes wrong, it must also explain why the skeptic's argument initially has such strong intuitive appeal. My solution does both. The argument from possibility AP2 goes wrong, because it is either unsound (due to a false premise as in MAP2 and EAP2<sub>x></sub>), or uninteresting and irrelevant to *fallible* knowledge (as in EAP2i), or invalid (due to equivocation as in EAP2<sub>x1</sub>), or question-begging (as in EAP2f). As a result, AP2 provides no good reason for thinking that I lack knowledge<sub>f</sub> that I have hands. Why then are so many people caught in AP2's skeptical grip, when first presented with the argument? The answer is that either: (i) having initially been drawn in by the  $\mu$ -possibility of the BIV hypothesis, they conflate  $\mu$ -possibility with  $e$ -possibility, thereby, in effect, illegitimately drawing an epistemic conclusion from purely metaphysical premises, or (ii) they recognize that the argument must be couched in terms of  $e$ -possibility, but they fail to notice the subtle equivocation between fallibilist and infallibilist senses of epistemic possibility identified in EAP2<sub>x1</sub>. Given the subtlety of each mistake, it is perfectly understandable that one find the unqualified AP2 initially threatening, indeed.

This way of undermining the skeptical problem posed by AP2 also allows us to resolve the paradox that EAP2 itself generates. The reason we are inclined to think that the BIV hypothesis is e-possible even though we *know* things incompatible with its e-possibility is because in making our BIV e-possibility assessment we are making an e-possibility<sub>i</sub> assessment, whereas in claiming to know that I have hands, we're making a knowledge<sub>f</sub> claim. Since we have very little, if any, *infallible* knowledge, nothing we *infallibly* know e-precludes the truth of the BIV hypothesis, and so, the BIV hypothesis is e-possible<sub>i</sub> for us. That is why we are initially seduced into accepting premise 1 of the ambiguous EAP2. The epistemological mistake that has been repeated for centuries and that most people make when first confronted with the e-possibility<sub>i</sub> of the BIV scenario is concluding, on that basis, that we have no knowledge<sub>f</sub>, which is just to fall prey to the equivocation identified in EAP2<sub>X1</sub>. The e-possibility<sub>i</sub> of the BIV hypothesis does prevent us from having knowledge<sub>f</sub> of the existence of the external world, as EAP2i shows. That's as it should be. But EAP2 is impotent when it comes to knowledge<sub>f</sub>, because there is no nonquestion-begging way to establish the e-possibility<sub>f</sub> of the BIV hypothesis.

This solution has the added virtue that it can explain why skeptical arguments typically lose their force for epistemologists. Once we realize that the skeptic cannot assert her major premise — that it's e-possible<sub>f</sub> that we are BIVs — without begging the question, we no longer find her arguments compelling. At first, we don't realize that she cannot rationally assert that premise without assuming the truth of her conclusion, but once we do, we find her arguments to be of absolutely no use in defending skepticism. This, I suspect, is why many epistemologists are no longer bothered by the very same skeptical arguments that at one time troubled them deeply.

#### 7. EPILOGUE: BACK TO THE VAT

At this point some of you may feel cheated. Haven't I pulled an end run around the skeptic (rather than confronting her) by substituting the argument from possibility for the BIV argument and arguing that the former is question-begging? I can imagine a skeptic objecting as follows:

You may be right that the argument from possibility begs the question, but what about the BIV argument with which we began? It seems immune to your question-begging charge, since it's valid, its second premise BIV1.2 follows from the closure

principle, and there is an independent argument for its first premise BIV1.1, an argument that makes no reference to BIV1's conclusion [BIV1.3].

I submit that the standard argument offered in support of BIV1.1 does presuppose the truth of BIV1.3. To see why, consider the following dialogue with me playing the role of the nonskeptic:

**Skeptic:** Do you know<sub>f</sub> that you have hands?

**Me:** Why yes, I think I do.

**Skeptic:** There's good reason to think you are mistaken.

**Me:** Really? I'd need a *very* good reason to give up such a commonsensical belief as the belief that I know<sub>f</sub> that I have hands. So tell me, what's the reason?

**Skeptic:** It's a very good reason, alright. It's based on the extremely plausible principle of epistemic closure according to which: If *S* knows<sub>f</sub> that *p*, and *S* knows<sub>f</sub> that *p* entails *q*, and *S* considers *q* in light of her knowledge<sub>f</sub> that *p* and that *p* entails *q*, and *S* bases her belief that *q* on this knowledge<sub>f</sub>, then *S* knows<sub>f</sub> that *q*. Surely, you accept the principle of epistemic closure, don't you?

**Me:** I certainly do.

**Skeptic:** Good, then I've got you. Here's the argument: (i) You don't know<sub>f</sub> you're not a handless BIV. The second premise derives from the closure principle, which you've just embraced: (ii) If you don't know<sub>f</sub> you're not a handless BIV, then you don't know<sub>f</sub> that you have hands. Therefore, (iii) you don't know<sub>f</sub> that you have hands.

**Me:** Hmm. Why should I accept (i)? It seems to me that you are presupposing that I don't know<sub>f</sub> I have hands in your first premise.

**Skeptic:** No, I'm not, I wouldn't expect you to accept (i) without an independent argument supporting it, but I have such an argument. It's the famous argument from possibility and runs as follows: (iv) Surely, it's *possible* that you are a handless BIV. (v) If it is possible that you are a handless BIV, then you don't know<sub>f</sub> that you're not a handless BIV. So, (vi) you don't know<sub>f</sub> you're not a handless BIV -- which is just my original premise (i).

**Me:** What kind of *possibility* are you appealing to in (iv)?

**Skeptic:** Epistemic possibility, of course.

**Me:** Fallibilistic or infallibilistic e-possibility?

**Skeptic:** Fallibilistic.

**Me:** Can we agree to the following account of fallibilistic e-possibility:

$D'_{3t}$   $p$  is e-possible<sub>t</sub> for  $S$  at  $t$  iff (i)  $S$  does not know<sub>t</sub> that  $\sim p$  at  $t$ , and (ii)  $S$  could not come to know<sub>t</sub> that  $\sim p$  at  $t$ , strictly on the basis of propositions  $S$  knows<sub>t</sub> at  $t$ ?

**Skeptic:** Yes, something like that seems correct.

**Me:** So, given  $D'_{3t}$ , your argument from possibility [our API] in support of (i) reduces to the following argument: (iv\*) I don't know<sub>t</sub> that I'm not a handless BIV, and I could not come to know<sub>t</sub> that I'm not a handless BIV, strictly on the basis of the propositions I currently know<sub>t</sub>. (v\*) If I don't know<sub>t</sub> that I'm not a handless BIV, and I could not come to know<sub>t</sub> that I'm not a handless BIV, strictly on the basis of the propositions I currently know<sub>t</sub>, then I don't know<sub>t</sub> that I'm not a handless BIV. Therefore, (vi) I don't know<sub>t</sub> that I'm not a handless BIV. Right?

**Skeptic:** Right.

**Me:** But then your argument from possibility fails to provide me with a nonquestion-begging reason to accept the first premise of your BIV argument. After all, the first conjunct of premise (iv\*) of your argument from possibility is "I don't know<sub>t</sub> that I'm not a handless BIV", which is precisely what your argument from possibility was supposed to establish.

**Skeptic:** Hmm.

**Me:** Worse still, you can only rationally assert the second conjunct of premise (iv\*) – that I could not come to know<sub>t</sub> that I am not a handless BIV, strictly on the basis of propositions I currently know<sub>t</sub> – if you make certain presuppositions about what I currently know<sub>t</sub>. In particular, you must assume that I don't currently know<sub>t</sub> that I have hands; for given your own commitment to closure, if I do currently know<sub>t</sub> that I have hands, then I could easily come to know<sub>t</sub> that I am not a handless BIV on the basis of my knowledge<sub>t</sub> that I have hands. Thus, in offering your argument from possibility in support of premise (i) of your BIV argument, you are presupposing the truth of the very conclusion your BIV argument was supposed to establish. You promised to provide a good argument for me to accept (iii) – the conclusion that I do not know<sub>t</sub> that I have hands – and now it turns out that the only way for you to defend the first premise of your

original DIV argument is to appeal to another argument which presupposes the truth of (iii) in its first premise, i.e. premise (iv). Such a pair of arguments is of no use in proving the truth of conclusion (iii). If there is a good argument for the skeptical conclusion "I don't know<sub>f</sub> that I have hands", you certainly have not produced it. Absent a good nonquestion-begging reason to believe that I don't know<sub>f</sub> that I have hands, it seems perfectly reasonable to retain the commonsense belief that I do know<sub>f</sub> that I have hands and, given closure, it's also reasonable to believe that I know<sub>f</sub> that I'm not a handless DIV.

#### ACKNOWLEDGEMENTS

An earlier version of this paper was presented at the "Contextualist Approaches to Epistemology" Conference in Mainz, Germany. I wish to thank my commentator, Gilbert Scharif, for his thoughtful commentary on my paper. I also wish to thank the conference participants for their helpful comments during the discussion period. Special thanks to Matthias Steup, Bruce Russell, John Carroll, Tomis Kapitan, James Dye, Elke Brendel, Christoph Jäger, and an anonymous referee for this journal for their extremely helpful, detailed comments on earlier drafts. The present version is much improved as a result.

#### NOTES

- <sup>1</sup> Skeptical hypotheses need to be sufficiently detailed alternative explanations of our experiences to give us pause. Lehrer details the Googol hypothesis as follows: "There are a group of creatures in another galaxy, call them Googols, whose intellectual capacity is  $10^{100}$  that of men, and who amuse themselves by sending out a peculiar kind of wave that affects our brain in such a way that our beliefs about the world are mostly incorrect. This form of error infects beliefs of every kind, but most of our beliefs, though erroneous, are nevertheless very nearly correct. This allows us to survive and manipulate our environments" (Lehrer 1971, p. 356).
- <sup>2</sup> Cohen endorses essentially the same adequacy constraint: "The burden of the fallibilist is to resolve these puzzles and paradoxes in a way that preserves the truth of our everyday knowledge attributions. But a *satisfying* resolution requires an explanation of why the paradox arises . . . an explanation of why we have the intuitions that saddle us with the paradox" (Cohen 1988, p. 94).
- <sup>3</sup> Epistemic closure is the thesis that knowledge is closed under known logical implication. Of course, as is widely recognized, we must be careful in how we



formulate the closure principle. For example, overly simplistic formulations, like Nozick's  $[K_s(p) \ \& \ K_s(p \rightarrow q)] \rightarrow K_s(q)$ , are false, for  $S$  may simply fail to put two and two together and thus not come to believe that  $q$ , in which case  $S$  will fail to know that  $q$ . [After presenting the above formulation, Nozick himself adds: "this principle counts on the person to draw the inference to  $q$ " (Nozick 1981, p. 205).] To avoid confusions that can be caused by such overly simplistic formulations, let us agree to understand the epistemic closure principle as follows: If  $S$  knows that  $p$ , and  $S$  knows that  $p$  entails  $q$ , and  $S$  considers  $q$  in light of her knowledge that  $p$  and that  $p$  entails  $q$ , and  $S$  bases her belief that  $q$  on this knowledge, then  $S$  knows that  $q$ .

- <sup>4</sup> DeRose refers to his version of contextualism as "attributor contextualism" (DeRose 1999, p. 190).
- <sup>5</sup> Lewis makes the point as follows: "I take it that the rule of accommodation can go both ways. But for some reason raising of standards goes more smoothly than lowering. If the standards have been high, and something is said that is true enough only under lowered standards, and nobody objects, then indeed the standards are shifted down. But what is said, although true enough under the lowered standards, may still seem imperfectly acceptable. Raising of standards, on the other hand, manages to seem commendable even when we know that it interferes with our conversational purpose. Because of this asymmetry, a player of language games who is so inclined may get away with it if he tries to raise the standards of precision as high as possible – so high, perhaps, that no material object whatever is hexagonal" (Lewis 1979, p. 352f). "We get the impression that the sceptic ... has the last word. Again this is because the rule of accommodation is not fully reversible. For some reason, I know not what, the boundary readily shifts outward if what is said requires it, but does not so readily shift inward if what is said requires that" (Lewis 1979, p. 355).
- <sup>6</sup> See DeRose (2000) for a fuller discussion of "lost knowledge".
- <sup>7</sup> For a formal demonstration that, in high standards contexts, one cannot know that one has low standards knowledge, see Brendel (2003).
- <sup>8</sup> DeRose acknowledges this very point: "Thus, on our solution, we do know, for instance, that we're not BIVs, according to ordinary low standards for knowledge. But, though (1) [of our BIV argument] is false when evaluated according to those ordinary low standards, we're able to explain its plausibility, as we've seen, by means of the fact that the high standards at which (1) is true are precisely the standards that an assertion or denial of it put into play. Since attempts to assert (1) are bound to result in truth, and attempts to deny it are destined to produce falsehood, it's no surprise that we find it so plausible" (DeRose 1995, p. 39f).

DeRose is so eager to explain the plausibility of premise (1) of the BIV argument that he doesn't seem to notice how counterintuitive it is to maintain that people have all sorts of unspeakable and unthinkable knowledge.

- <sup>9</sup> Some philosophers regard *arguments from possibility* as the most fundamental skeptical arguments. See, e.g. Vogel (forthcoming), where he writes: "The argument which supports skepticism is one of the most famous in the history of philosophy. It turns on the possibility that we might be victims of some kind of massive sensory deception." Also see: (a) Stroud (1984, chapter 1), where he develops the Cartesian argument from possibility at length; and (b) Nozick's discussion of "Skeptical Possibilities" (Nozick 1981, pp. 198–204).

- <sup>10</sup> For present purposes, let us stipulatively agree to the following: A proposition  $p$  is  $\mu$ -possible iff there is a  $\mu$ -possible world where  $p$  is true. A proposition  $p$  is  $\mu$ -impossible iff there are no  $\mu$ -possible worlds where  $p$  is true.
- <sup>11</sup> Obviously, condition (ii) needs unpacking. Here is what is intended by condition (ii):
- D<sub>4</sub>  $S$  could come to know that  $\sim p$  at  $t$ , strictly on the basis of the propositions  $S$  knows at  $t$ , iff either (1) one or more of the propositions  $S$  knows at  $t$  self-evidently entail that  $\sim p$  for  $S$  (such that it is within  $S$ 's cognitive capacity at  $t$  to grasp that entailment immediately at  $t$ ); or (2)  $\sim p$  is true, one or more of the propositions  $S$  knows at  $t$  provide an adequate justificatory basis for believing that  $\sim p$ , and it is within  $S$ 's cognitive capacity at  $t$  to see that these propositions justify her in believing that  $\sim p$ .
- <sup>12</sup> It's worth noting that D<sub>3</sub> entails that no true proposition is ever e-impossible for  $S$ , because if  $p$  is true, then  $\sim p$  is false, and since  $\sim p$  is false,  $S$  couldn't know that  $\sim p$ . At first blush, it might seem surprising that no true proposition is ever e-impossible. Nevertheless, the result is correct. How, after all, could something  $S$  knows to be true e-preclude the truth of another true proposition?
- <sup>13</sup> D<sub>3</sub> can be stated more explicitly as follows:
- D<sub>3</sub><sup>c</sup>  $p$  is e-possible for  $S$  at  $t$  iff (i)  $S$  does not know that  $\sim p$  at  $t$ ; (ii) if one or more of the propositions  $S$  knows at  $t$  entail that  $\sim p$ , then it is not within  $S$ 's cognitive capacity at  $t$  to grasp that entailment; and (iii) if  $\sim p$  is true and if one or more of the propositions  $S$  knows at  $t$  provide an adequate justificatory basis for believing that  $\sim p$ , then it is not within  $S$ 's cognitive capacity at  $t$  to see that these propositions justify her in believing that  $\sim p$ .
- <sup>14</sup> D<sub>3</sub> also entails the right result regarding Simone. I know that Simone is in the bedroom, because I reliably, justifiably, and truly believe that she is lying in bed next to me. I do not know that *Simone is not in the kitchen*, because I haven't considered that proposition. However, some of the propositions I know – (a) that *Simone is in the bedroom*, and (b) that *the bedroom and the kitchen are distinct rooms in my palatial estate* – self-evidently entail that *Simone is not in the kitchen*, and I am quite capable of grasping that entailment. So, it is not e-possible for me that *Simone is in the kitchen*, because, even though I don't know that she is not in the kitchen, some of the things I do know at the time, viz. (a) and (b), obviously entail that she is not in the kitchen.
- <sup>15</sup> Where:  $\diamond$  = It is possible that ...;  $K$  = I know that ...;  $b$  = I am a handless BIV;  $h$  = I have hands.
- <sup>16</sup> Where:  $\diamond_{\mu}$  = It is metaphysically possible that ...
- <sup>17</sup> To think otherwise is the epistemic equivalent of the naturalistic fallacy.
- <sup>18</sup> I am not begging the question against the skeptic here. I am simply making a conceptual point that from the mere fact that Bob knows that it is  $\mu$ -possible that Jim is in New Orleans, it does not follow that Bob does not know that Jim is not in New Orleans, for as we have seen before, Bob may know that the  $\mu$ -possibility in question is not actual.
- <sup>19</sup> Where:  $\diamond_e$  = It is e-possible that ...
- <sup>20</sup> Many students, of course, when first presented with an argument like RAP2, feel compelled to accept the conclusion, because they cannot find anything wrong with the argument.

- <sup>21</sup> It is not idiosyncratic of  $D_3^e$  that it entails that  $S$ 's not knowing that  $\sim p$  is a necessary condition for  $p$ 's being e-possible for  $S$ . In fact, virtually every purported account of e-possibility in the literature implies that  $S$ 's not knowing that  $\sim p$  is necessary for  $p$ 's being e-possible for  $S$ .
- <sup>22</sup> DeRose defends a similar principle in DeRose 1991, pp. 596-601.
- <sup>23</sup> Infallibilistic intuitions clearly underlie Saul Kripke's *a priori* Cartesian certainty account of e-possibility. See Kripke (1980, p. 143, fn. 72), where he proposes the following:  $p$  is e-possible for  $S$  iff  $S$ 's evidence does not justify *a priori* Cartesian certainty that  $\sim p$ .
- <sup>24</sup> Knowledge<sub>e</sub> and knowledge<sub>eT</sub> are distinguished as follows: Knowledge<sub>e</sub> requires infallible justification, i.e. justification that entails that which it justifies.  $S$  is justified, in believing that  $p$  [ $J_1p$ ] only if  $p$ . Knowledge<sub>eT</sub> can be analyzed as follows:

$$(K_1) \quad K_1p \equiv (p \ \& \ Bp \ \& \ J_1p).$$

According to fallibilism, the kind of justification needed to convert true belief to knowledge<sub>e</sub> must only make probable, but need not entail, that for which it is justification. As a result, fallibilism entails:  $\Diamond(J_1p \ \& \ \sim p)$ . This possibility and the closure principle with respect to justification together entail numerous Gettier-possibilities, including:

$$\Diamond[Bp \ \& \ J_1p \ \& \ J_1(p \ \rightarrow q) \ \& \ B(p \ \rightarrow q) \ \& \ J_1q \ \& \ Bq \ \& \ q \ \& \ \sim p \ \& \ \sim K_1q]$$

The latter possibility obtains when, as Gettier illustrated,  $S$  has a justified<sub>e</sub> true belief that  $q$  which falls short of knowledge<sub>eT</sub> because  $S$ 's justification<sub>e</sub> for  $q$  (to wit,  $Bp \ \& \ J_1p \ \& \ J_1(p \ \rightarrow q) \ \& \ B(p \ \rightarrow q)$ ) fails to be appropriately connected to  $q$ 's truth and thus is defective (Gettier 1963, pp.121-123). Since fallibilism entails these possibilities, a fourth condition must be added to the traditional analysis of knowledge to rule out Gettier cases as instances of knowledge. For our purposes, the following condition will suffice:  $S$  is not Gettierized with respect to  $p$  [ $\sim Gp$ ]. Accordingly, we can analyze knowledge<sub>eT</sub> as follows:

$$(K_2) \quad K_2p \equiv (p \ \& \ Bp \ \& \ J_2p \ \& \ \sim Gp).$$

- <sup>25</sup>  $D_3^e$  should be revised accordingly:

$D_3^e$   $p$  is e-possible<sub>e</sub> for  $S$  at  $t$  iff (i)  $S$  does not know<sub>e</sub> that  $\sim p$  at  $t$ ; (ii) if one or more of the propositions  $S$  knows<sub>e</sub> at  $t$  entail that  $\sim p$ , then it is not within  $S$ 's cognitive capacity at  $t$  to grasp that entailment; and (iii) if  $\sim p$  is true and if one or more of the propositions  $S$  knows<sub>e</sub> at  $t$  provide an adequate justificatory<sub>e</sub> basis for believing that  $\sim p$ , then it is not within  $S$ 's cognitive capacity at  $t$  to see that these propositions justify<sub>e</sub> her in believing that  $\sim p$ .

$D_3^e$   $p$  is e-possible<sub>e</sub> for  $S$  at  $t$  iff (i)  $S$  does not know<sub>e</sub> that  $\sim p$  at  $t$ , and (ii) if one or more of the propositions  $S$  knows<sub>e</sub> at  $t$  entail that  $\sim p$ , then it is not within  $S$ 's cognitive capacity at  $t$  to grasp that entailment.

An infallibilistic version of condition (iii) of  $D_3^e$  is redundant where e-possibility, is concerned.

<sup>26</sup> Where:  $\Diamond_c$  = It is c-possible<sub>c</sub> that ...;  $K_c$  = I know<sub>c</sub> that ...

<sup>27</sup> Where:  $\Diamond_c$  = It is c-possible<sub>c</sub> that ...;  $K_c$  = I know<sub>c</sub> that ...

<sup>28</sup> It is precisely such reasoning that inclines us to accept premise 1 of the original ambiguous EAP2.

<sup>29</sup> We might make the point as follows. Because EAP2f is valid, so is the following argument:

2.  $\Diamond_c b \rightarrow \Diamond_c \sim h$
  3.  $\Diamond_c \sim h \rightarrow \sim K_c h$
  - $\sim$  4.  $K_c h$  (i.e.  $\sim \sim K_c h$ )
- $\therefore \sim$  1.  $\sim \Diamond_c b$

The above argument demonstrates that the falsity of EAP2f's conclusion entails the falsity of EAP2f's first premise, since, as we have seen, EAP2f 2 and 3 are true.

<sup>30</sup> Peter Klein makes a similar point with respect to skeptical arguments predicated on the closure principle. He claims that such arguments "virtually beg the question" because one of the premises in closure-based skeptical arguments can only be supported by a subargument that employs the conclusion of the main skeptical argument as a premise. See Klein (1995). While Klein properly diagnoses one way skeptical arguments can go wrong, he does not explain the source of their intuitive appeal, nor does he acknowledge the role equivocation plays in motivating skepticism.

#### REFERENCES

- Brendel, E.: 2003, 'Was Kontextualisten Nicht Wissen', *Deutsche Zeitschrift für Philosophie* 51(6), 1015–1032.
- Cohen, S.: 1988, 'How to Be a Fallibilist', *Philosophical Perspectives* 2 (Issue Epistemology), 91–123.
- DeRose, K.: 1991, 'Epistemic Possibilities', *The Philosophical Review* 100(4), 581–605.
- DeRose, K.: 1995, 'Solving the Skeptical Problem', *The Philosophical Review* 104(1), 17–52.
- DeRose, K.: 1999, 'Contextualism: An Explanation and Defense', in J. Greco and E. Sosa (eds), *The Blackwell Guide to Epistemology*, Blackwell, Malden, MA, Oxford, pp. 187–205.
- DeRose, K.: 2000, 'Now You Know It, Now You Don't', *Proceedings of the Twentieth World Congress of Philosophy* Vol. 5, Epistemology, Philosophy Documentation Center, Bowling Green, Ohio, pp. 91–106.
- Gettier, E.: 1963, 'Is Justified True Belief Knowledge?', *Analysis* 23, 121–123.
- Hacking, I.: 1975, 'All Kinds of Possibility', *The Philosophical Review* 84, 321–337.
- Hume, D.: 1739, *A Treatise of Human Nature*, Second Selby-Bigge's edition (1978), Clarendon, Oxford.
- Klein, P.: 1995, 'Skepticism and Closure: Why the Evil Genius Argument Fails', *Philosophical Topics* 23, 213–236.
- Kripke, S.: 1980, *Naming and Necessity*, Harvard University Press, Cambridge, MA.

- Lehrer, K.: 1971, 'Why Not Scepticism?', *The Philosophical Forum* 2.3, pp. 283–298 (reprinted in G. Pappas and M. Swain (eds), *Essays on Knowledge and Justification*, Cornell University Press, Ithaca and London, pp. 346–363, pages references to the Pappas and Swain volume.)
- Lewis, D.: 1979, 'Scorekeeping in a Language Game', *Journal of Philosophical Logic*, 8, 339–359.
- Lewis, D.: 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* 74(4), 549–567.
- Nozick, R.: 1981, *Philosophical Explanations*, Harvard University Press, Cambridge, MA.
- Stroud, B.: 1984, *The Significance of Philosophical Skepticism*, Clarendon Press, Oxford.
- Vogel, J.: forthcoming, 'The Refutation of Skepticism', in E. Sosa and M. Stoup (eds), *Contemporary Debates in Epistemology*, Blackwell, Malden, MA, Oxford.

Department of Philosophy  
Northern Illinois University  
De Kalb, IL 60115  
USA