Constraints on Skeptical Hypotheses

Forthcoming in *The Philosophical Quarterly*

James R. Beebe (University at Buffalo)

In this article I examine the constraints that skeptical hypotheses must satisfy in order to be used to raise significant skeptical challenges. I argue that skeptical hypotheses do not have to be logically, metaphysically or epistemically possible. They only need to depict scenarios that are subjectively indistinguishable from the actual world and must provide some indication of how subjects can believe what they do while failing to have knowledge. I also argue that skeptical challenges can be raised against *a priori* beliefs, even if those beliefs are necessarily true. In this way I hope to broaden our conception of the legitimate kinds of skeptical challenges that can be raised.

Skeptical hypotheses depict situations that are subjectively indistinguishable from what we take our normal circumstances to be but in which we fail to have knowledge. There are several constraints that skeptical hypotheses must satisfy in order to underwrite effective skeptical challenges. Many widespread views about these constraints, however, are mistaken. For example, it is widely but incorrectly believed that skeptical hypotheses must describe scenarios in which subjects' beliefs are false and that skeptical hypotheses must be logically or metaphysically possible. In this article I use a series of thought experiments to probe the set of requirements skeptical hypotheses must satisfy and argue that effective skeptical hypotheses do not have to be logically, metaphysically or even epistemically possible. Subjective indistinguishability (understood as a form of subjective possibility) is all that is required. I also show that neither the necessary truth nor the *a priori* status of a belief can render it immune to skeptical attack.

I begin in section I by arguing that skeptical hypotheses need not be incompatible with what subjects believe but that they must provide some indication of how subjects can believe what they do while failing to have knowledge. In section II I argue that skeptical hypotheses can be used to raise epistemological challenges to beliefs that are necessarily true. In the following two sections I argue that skeptical hypotheses can be used to challenge instances of putative *a priori* knowledge (section III), even if those hypotheses are logically or metaphysically impossible (section IV). In section V I argue against the view that skeptical hypotheses must be epistemically possible, and in section VI I articulate a subjective indistinguishability constraint on skeptical hypotheses.

It is important to note that I am primarily concerned with the question of which constraints skeptical hypotheses must satisfy in order to pose significant skeptical challenges. I do not attempt to elucidate the proper constraints on successful replies to skepticism or to determine whether the various skeptical challenges I discuss can be effectively answered. My goal is to illuminate the nature of those challenges themselves.

I.

The most familiar skeptical arguments in the contemporary literature rely upon a closure principle for knowledge such as the following:

(CP) If *S* knows that *p* and *S* knows that *p* entails *q*, then *S* knows (or is in a position to know) that *q*.

Where *O* is a proposition we ordinarily take ourselves to know and *SK* is an appropriately chosen skeptical hypothesis, the following is perhaps the most commonly encountered form of skeptical argument:

(1.1) If I know that *O*, then I know that not-*SK*.

(1.2) I don't know that not-*SK*.

(1.3) Therefore, I don't know that *O*.

Another well-known skeptical argument appeals to considerations of underdetermination:

(2.1) If my evidence for believing that *O* does not favor *O* over some hypothesis, *SK*, which I know to be incompatible with *O*, then my evidence does not justify me in believing *O*.

(2.2) My evidence for believing that *O* does not favor *O* over *SK*.

(2.3) Therefore, I'm not justified in believing that *O*.

(2.4) Therefore, I don't know that *O*.

The first premise of this argument is based upon an underdetermination principle such as the following:

(UP) If *S*'s evidence for believing that *p* does not favor *p* over some hypothesis *q* which *S* knows to be incompatible with *p*, then *S*'s evidence does not justify *S* in believing *p*.[1]

Although many epistemologists (e.g., Brueckner 1994; Cohen 1998; DeRose 1999b) claim that closure- and underdetermination-based skeptical arguments capture the heart of the historic skeptical challenge, (CP) and (UP) can only be used in conjunction with skeptical hypotheses that satisfy the following constraint:

(SH1) In order for a skeptical hypothesis, *SK*, to raise a significant skeptical challenge to *S*'s putative knowledge that *O*, *SK* must be incompatible with *O*.

However, it should be well known that dreaming hypotheses can depict situations that are compatible with what we ordinarily believe and yet can underwrite effective skeptical challenges to our knowledge of the external world. G. E. Moore (1959, 245) vividly portrays this point with the following anecdote:

> But, on the other hand, from the hypothesis that I am dreaming, it certainly would not follow that I am *not* standing up; for it is certainly logically possible that a man should be fast asleep and dreaming, while he is standing up and not lying down. It is therefore logically possible that I should both be standing up and at the same time dreaming that I am; just as the story, about a well-known Duke of Devonshire, that he once dreamt that he was speaking in the House of Lords

and, when he woke up, found that he *was* speaking in the House of Lords, is certainly logically possible.

Since a subject cannot know that *O* on the basis of dreaming that *O*, the skeptic can argue that the subject needs to be able to rule out the dreaming skeptical hypothesis in order to know that *O*. Compare the following statements:

(3.1) Anyone who recognizes the incompatibility between having hands and being a brain in a vat must be in a position to know that she is not a brain in a vat in order to know that she has hands.

(3.2) Anyone who recognizes the incompatibility between knowing that one is standing and dreaming that one is standing must be in a position to know that she is not dreaming in order to know that she is standing.

Descartes and Moore certainly would have thought (3.2) to be no less plausible than (3.1). If this assessment is correct, the following arguments should be of comparable strength and plausibility (assuming in each case that the subject in question recognizes the relevant incompatibility):

(4.1) If I know that I have hands, then I know that I'm not a brain in a vat.

(4.2) I don't know that I'm not a brain in a vat.

(4.3) Therefore, I don't know that I have hands.

(5.1) If I know that I'm standing, then I know that I'm not merely dreaming that I'm standing.

(5.2) I don't know that I'm not merely dreaming that I'm standing.

(5.3) Therefore, I don't know that I'm standing.

Thus, a dreaming skeptical hypothesis can raise a significant challenge to a subject's putative knowledge that *O* without depicting a situation that is incompatible with *O*. (SH1), therefore, is false.[2]

The following constraint on skeptical hypotheses is also incorrect:

(SH2) In order for a skeptical hypothesis, *SK*, to raise a significant skeptical challenge to *S*'s putative knowledge that *O*, it is sufficient that *SK* be incompatible with *O*.

Consider Fred, who believes that the animal standing before him in the pen at the zoo is a zebra. The following propositions are all incompatible with what Fred believes:

(6.1) The animal in the pen is a lion.

(6.2) The animal in the pen is not a zebra.

(6.3) The animal in the pen is a mule cleverly disguised to look like a zebra.

Clearly, however, (6.1) and (6.2) do not represent skeptical hypotheses.[3] All three propositions satisfy (SH1), but only (6.3) satisfies the following, broadly explanatory constraint:

(SH3) In order for a skeptical hypothesis, *SK*, to raise a significant skeptical challenge to *S*'s putative knowledge that *O*, *SK* must indicate how *S* could believe that *O* on the basis of *S*'s evidence and yet fail to know that *O*.[4]

(6.1) and (6.2) show how Fred could fail to know that the animal in the pen is a zebra—viz., by believing something that is false—but they fail to provide any indication of how Fred could believe that the zebra-looking animal before him is a lion or otherwise not a zebra.[5] (SH2), then, is false, while (SH3) seems true.

II.

An erroneous constraint on skeptical hypotheses that seems to have broad appeal is the following:

(SH4) In order for a skeptical hypothesis, *SK*, to raise a significant skeptical challenge to *S*'s putative knowledge that *O*, it must be logically or metaphysically possible for *O* to be false.

The falsity of (SH4) can be revealed by noting that whether an effective skeptical challenge to religious belief (or unbelief) can be raised seems to have nothing to do with whether or not a divine being actually

exists. Theists, for example, believe that God exists, and atheists believe that God does not exist (where 'God' in each case denotes a necessarily existent divine being). One of these beliefs is necessarily true, while the other is necessarily false. According to (SH4), skeptical challenges can only be raised against one of these beliefs—the one that is necessarily false. But that is absurd. Suppose that God exists. Would this mean that no skeptical challenge to belief in God could ever be raised? Surely not.

Freud (1927/1961, 30) offered the following, undermining explanation of religious belief, which he took to be compatible with God's existence:

> These [religious beliefs], which are given out as teachings, are not precipitates of experience or end-results of thinking: they are illusions, fulfillments of the oldest, strongest and most urgent wishes of mankind. The secret of their strength lies in the strength of those wishes. As we already know, the terrifying impressions of helplessness in childhood aroused the need for protection—for protection through love—which was provided by the father; and the recognition that this helplessness lasts throughout life made it necessary to cling to the existence of a father, but this time a more powerful one. Thus the benevolent rule of a divine Providence allays our fear of the dangers of life; the establishment of a moral world-order ensures the fulfillment of the demands of justice, which have so often remained unfulfilled in human civilization; and the prolongation of earthly existence in a future life provides the local and temporal framework in which these wish-fulfillments shall take place.

According to Freud, then, belief in the existence of God arises from a psychological mechanism aimed at wish-fulfillment. In Freud's terminology, belief in God is an *illusion* rather than a *delusion* because it is not necessarily false. Although a proposition that is believed as the result of wish-fulfillment can be true, it cannot be known on that basis. Freud's religious skeptical hypothesis is thus analogous to Descartes's dreaming skeptical hypothesis.

Accordingly, a religious skeptic could offer the following argument against a theist's putative knowledge of the existence of God, even if that belief is necessarily true:

(7.1) If you know that God exists, then you know that your belief in God is not produced by a psychological mechanism aimed at wish-fulfillment.

(7.2) You don't know that your belief in God is not produced by a psychological mechanism aimed at wish-fulfillment.

(7.3) Therefore, you don't know that God exists.

(We are again assuming that the subject recognizes the incompatibility between knowing that God exists and believing that God exists on the basis of wish-fulfillment.) Commenting on the skeptical challenge raised by Freudian explanations of religious belief, Alvin Plantinga (2000, 195) writes:

> [T]he beauty of Freudian explanations is that the postulated mechanisms all operate unconsciously, unavailable to inspection. The claim is that you subconsciously recognize the miserable and frightening condition we human beings face, subconsciously see that the alternatives are paralyzing despair or belief in God, and subconsciously opt for the latter. Even after careful introspection and reflection, you can't see that the proffered explanation is true: that fact won't be taken as even the slightest reason for doubting the explanation.

Thus, even if a necessary being were to exist, this fact alone would not insulate religious belief against skeptical attack.

It is rarely appreciated that religious skeptical challenges can be run in the other direction. Theists, that is, can offer undermining explanations of religious unbelief such as the following. Drawing inspiration from Thomas Aquinas and John Calvin, Plantinga (2000) claims that human beings have a *sensus divinitatus*—an innate cognitive faculty that, when functioning properly, produces in us the belief that God exists. Plantinga (2000, 184) believes that the *sensus divinitatus* has been "compromised, weakened, reduced, smothered, overlaid, or impeded by sin and its consequences" ever since the Fall of Adam and Eve.

Because it has been damaged and corrupted by sin in this way, it may be partly or wholly disabled. Plantinga (2000, 184) writes:

> There is such a thing as cognitive disease; there is blindness, deafness, inability to tell right from wrong, insanity; and there are analogues of these conditions with respect to the operation of the *sensus divinitatus*.... [A]ccording to the model, it is really the *unbeliever* who displays epistemic malfunction; failing to believe in God is a result of some kind of dysfunction on the *sensus divinitatus*.

Thus, even if atheism were necessarily true, Plantinga's (2000) account of the *sensus divinitatus* and the cognitive consequences of sin could function as a skeptical hypothesis that challenged the atheist's knowledge of this fact. Like the ordinary skeptic, theists could argue that in order to know that atheism is true, atheists must be in a position to know that Plantinga's account is false but that they cannot know such a thing. Consequently, the necessary truth of a belief cannot neutralize all skeptical challenges to that belief. (SH4), therefore, is false.

It might be helpful at this point to remind ourselves that the following constraint on skeptical hypotheses is also incorrect:

> (SH5) In order for a skeptical hypothesis, *SK*, to raise a significant skeptical challenge to *S*'s putative knowledge that *O*, it must be plausible to believe that *SK* is true.

Brain in a vat and evil demon hypotheses do not raise effective skeptical challenges because anyone thinks it is plausible to believe we really are brains in vats or deceived by an evil demon. This means that one cannot argue that Plantinga's hypothesis fails to raise a legitimate skeptical challenge to religious unbelief on the grounds that it is implausible to suppose that we have a *sensus divinitatus*. The degree to which a skeptical hypothesis is plausible may play a role in determining how difficult it can be to dismiss skeptical challenges that are based upon it, but the implausibility of the hypothesis itself cannot keep those skeptical challenges from being raised in the first place.

III.

All of the skeptical hypotheses we have considered thus far—and indeed practically all of the skeptical hypotheses encountered in the contemporary literature—are used to challenge instances of allegedly *a posteriori* knowledge. Many philosophers are skeptical [sic] of the possibility of there being skeptical challenges to *a priori* knowledge. Matthias Steup (2005, 10-11), for example, writes:

> It is generally agreed that PAPs [i.e., putatively *a priori* propositions] are knowable. There is skepticism about knowledge of the external world, other minds, and the past. Skepticism about PAPs, however, is rarely pursued. Indeed, considering that knowledge of PAPs includes knowledge of the laws of logic, and more specifically, knowledge of an argument's validity, it is hard to see how a skeptical argument for anything could get off the ground without the prior assumption that knowledge of PAPs is indeed possible.

The following constraint on skeptical hypotheses appears to be widely endorsed, if only because of the almost total neglect of skeptical hypotheses that fail to satisfy it:

> (SH6) In order for a skeptical hypothesis, *SK*, to raise a significant skeptical challenge to *S*'s putative knowledge that *O*, *S*'s putative knowledge must be *a posteriori*.[6]

Like many other alleged constraints on skeptical hypotheses, however, (SH6) can be shown to be false with a modest amount of critical reflection.

Consider the following skeptical hypothesis: A bumbling evil demon is intent upon deceiving his subjects about certain *a priori* matters. He notes that there seems to be a distinct kind of phenomenology associated with rational intuitions—i.e., mental episodes in which *a priori* propositions intellectually seem to be true. George Bealer (2004, 12) describes rational intuitions in the following way:

> By intuitions we mean *seemings:* for you to have an intuition that *p* is just for it to *seem* to you that *p*.... For example, when you first consider one of de Morgan's laws, often it neither seems true nor seems false; after a moment's reflection, however, something happens: it now just seems

> true. This kind of seeming is *intellectual*, not experiential—sensory, introspective, imaginative. Intuition is different from belief: you can believe things that you do not intuit (e.g., that Fribourg is in Switzerland), and you can intuit things that you do not believe (e.g., the axioms of naive set theory). The experiential parallel is that you can believe things that do not appear (seem sensorily) to be so, and things can seem sensorily in ways you do not believe them to be (as with the Müller-Lyer arrows).

Plantinga (1993, 104) claims there is a distinct kind of phenomenology associated with rational intuitions—a "feeling of rightness or correctness":

> [C]onsidering or entertaining *If all men are mortal and Socrates is a man, then Socrates is mortal* feels different, somehow, from considering, say, *If all men are mortal and Lassie is mortal, then Lassie is a man*. The one belief seems right, compelling, acceptable; the other seems wrong, off-putting, and eminently rejectable; and this difference in experience is surely connected with our accepting the one and rejecting the other.

Suppose the bumbling evil demon attempts to deceive his subjects by switching the two types of experiences Plantinga describes, making a consideration of affirming the consequent be accompanied by a feeling that it is "right, compelling, acceptable" and *modus ponens* seem "wrong, off-putting and eminently rejectable." However, because the evil demon is not very practiced in the art of deception, he mistakenly makes affirming the consequent seem wrong and *modus ponens* seem right. If his victims were to base their beliefs in the merits of *modus ponens* and affirming the consequent on the intellectual seemings provided to them by the evil demon, their beliefs would not count as knowledge, however true they might be.[7] (The depicted scenario is thus an *a priori* Gettier case.[8])

The intuitive experiences of subjects in the foregoing scenario are subjectively indistinguishable from those had by subjects in "normal" situations (where this means their intuitive experiences arise from a proper *a priori* grasp of the propositions in question). Yet subjects in the latter situation have knowledge,

whereas subjects in the former do not. The following skeptical argument can be based upon this *a priori* skeptical hypothesis:

(8.1) If I know that *modus ponens* is correct, then I know that my belief that *modus ponens* is correct is not based on faux intuitive experiences that are the work of a bumbling evil demon.

(8.2) I don't know that my belief that *modus ponens* is correct is not based on faux intuitive experiences that are the work of a bumbling evil demon.

(8.3) Therefore, I don't know that *modus ponens* is correct.

The *a priori* skeptic can note the strong parallel between this argument and the dreaming skeptical argument from (5.1) to (5.3). Since the deceptive work of the bumbling evil demon would prevent me from having *a priori* knowledge, it seems that I must know the falsity of this skeptical hypothesis in the same way that I must know the falsity of the dreaming skeptical hypothesis in order to have knowledge of the external world. Furthermore, it can be difficult to see how I could know the falsity of either hypothesis, since if they were true my experiences would be exactly as they are. Knowledge of *a priori* necessities thus seems as vulnerable to skeptical attack as more traditional targets. (SH6), then, is false.

IV.

All of the skeptical hypotheses described above satisfy the following constraint:

(SH7) In order for a skeptical hypothesis, *SK*, to raise a significant skeptical challenge to *S*'s putative knowledge that *O*, it must be logically or metaphysically possible for *SK* to be true.

It is clearly possible for all religious beliefs to be formed on the basis of wish-fulfillment and for a variety of *a priori* beliefs to be formed on the basis of spurious intellectual seemings. The question I now want to

consider is whether effective skeptical hypotheses can be logically or metaphysically impossible. I will argue that they can be.

I begin by noting that Descartes believed that skeptical hypotheses involving apparent impossibilities were not beyond the pale of serious philosophical consideration. For example, in the first *Meditation* Descartes considers the possibility that an all-powerful being might be deceiving him about elementary propositions that seem to be grasped by reason alone:

> What is more, since I sometimes believe that others go astray in cases where they think they have the most perfect knowledge, may I not similarly go wrong every time I add two and three or count the sides of a square, or in some even simpler matter, if that is imaginable? (AT VII:21; CSM II:14)

In the *Third Meditation* Descartes reflects upon the possibility that "some God could have given me a nature such that I was deceived even in matters which seemed most evident":

> I cannot but admit that it would be easy for him, if he so desired, to bring it about that I go wrong even in those matters which I think I see clearly with my mind's eye. (AT VII: 36; CSM II: 25)

Descartes takes the force of this skeptical threat to be intensified when he considers that his origins and nature might be the result of mere chance rather than the handiwork of a perfect, omnipotent God:

> According to their supposition, then, I have arrived at my present state by fate or chance or a continuous chain of events, or by some other means; yet since deception and error seem to be imperfections, the less powerful they make my original cause, the more likely it is that I am so imperfect as to be deceived all the time. (AT VII: 21; CSM II:14)

Descartes's skeptical hypothesis attempts to cast doubt upon our ability to know putatively necessary truths on the basis of their intellectually seeming to be true by describing a situation in which we would have the same intuitions without those propositions being true.

An analogous but perhaps more effective skeptical hypothesis is suggested by Wittgenstein's reflections on logical necessity, at least as those reflections have been interpreted by Barry Stroud (1965). According to Stroud, Wittgenstein tries to steer a middle course between: (i) full-blooded conventionalism, which takes the necessity of any statement to consist in our having expressly decided to treat that statement as unassailable, and (ii) a Platonic realism, which locates the source of logical necessity in mind-independent facts. In Stroud's opinion, Wittgenstein agrees with realists that we can have no clear understanding of what it would mean for the apparently necessary truths of mathematics and logic to be false. Yet Wittgenstein also agrees with conventionalists that our ways of inferring, counting, calculating and so on are not the only possible ones. In his *Remarks on the Foundations of Mathematics*, Wittgenstein suggests that the following mathematical practices might represent genuine alternatives to our own:

(9.1) Following the rule "+ 2" by constructing the series "2, 4, 6,..., 996, 998, 1000, 1004, 1008,...."

(9.2) Agreeing that *modus ponens* is deductively valid, yet failing to agree that *q* follows from *p* and if *p* then *q*.

(9.3) Measuring with rulers that expand to an extraordinary extent when slightly heated.

(9.4) Dividing by (*n*-*n*) and not being bothered by the results.

(9.5) Selling wood according to the area covered by a pile of wood.

(9.6) Selling wood at a price equal to the labor of felling the timber, measured by the age and strength of the woodsman.

Wittgenstein denies that we can know that the reason such alternatives are unimaginable to us is that they lead to logical contradictions. They may not be real possibilities for creatures like us and they may not be fully intelligible to us, but Wittgenstein wants to insist that they are nonetheless possibilities in some sense.

Describing his reflections on our mathematical practices, Wittgenstein writes:

> What we are supplying are really remarks on the natural history of man; not curiosities however, but rather observations on facts which no one has doubted, and which have only gone unremarked because they are always before our eyes. (1956, I, §141)
>
> I am not saying: if such-and-such facts of nature were different people would have different concepts (in the sense of a hypothesis). But: if anyone believes that certain concepts are absolutely the correct ones, and that having different ones would mean not realizing something that we realize—then let him imagine certain very general facts of nature to be different from what we are used to, and the formation of concepts different from the usual ones will become intelligible to him. (1953, IIxii)

Thus, Wittgenstein believes that if the "natural history" of our species had gone differently, we might have had different concepts and found different things to be conceivable, inconceivable or natural. Stroud (1965, 513) writes:

> It is in that sense a contingent fact that calculating, inferring, and so forth, are carried out in the ways that they are—just as it is a contingent fact that there is such a thing as calculating or inferring at all. But we can understand and acknowledge the contingency of this fact, and hence the possibility of different ways of calculating, and so forth, without understanding what those different ways might have been.

Instead of asking readers to imagine what it would be like for our basic logical and mathematical beliefs to be false, Wittgenstein asks his readers to consider something like the following:

> (9.7) Creatures like us might have compelling intuitions about what constitutes correct calculating, reasoning or measuring, even if those intuitions have no essential connection to the facts (if any) about what correct calculating, reasoning and measuring consist in.

(9.7) seems to represent a genuine possibility for us (even if the sense in which it is possible requires some explication[9]). Furthermore, it seems the *a priori* skeptic can ask how we know that this possibility does not

represent our actual situation. Since there is an incompatibility between my having genuine knowledge of necessary truths and the obtaining of the possibility described in (9.7), it seems I must be in a position to rule out this possibility if I am to have any knowledge of logical or mathematical truths. Wittgenstein's account of logical necessity can thus be co-opted to serve as a skeptical hypothesis, even though it was not intended to serve as one.

Descartes and Wittgenstein, then, each describe scenarios in which mental episodes of seeming to see that certain *a priori* propositions are self-evidently true do not indicate the truth of those propositions. Like skeptical hypotheses about the external world, these *a priori* skeptical hypotheses show how it is possible for certain classes of appearances to fail to reflect reality. Let a 'DW' be any subject whose *a priori* beliefs are massively and constantly in error due to the sorts of circumstances described by Descartes and Wittgenstein. The following skeptical argument can be constructed on the basis of the foregoing hypotheses:

(10.1) If I know that 2 + 3 = 5, then I know that I'm not a DW.[10]

(10.2) I don't know that I'm not a DW.

(10.3) Therefore, I don't know that 2 + 3 = 5.

Call the form of skepticism supported by this argument '*a priori* skepticism.' The *a priori* skeptic can note that if I were a DW, I would falsely believe I wasn't one, and could argue on this basis that I fail to know I'm not a DW, even if my belief is correct.[11] Furthermore, since my intuitive evidence or experience would be exactly what it is now if I were a DW, the skeptic could argue that no appeal to that evidence or experience could suffice to show that I live in a "normal" world (i.e., a world where I really do grasps necessary truths *a priori*) rather than a DW world. (SH7), then, is false.

Note that one cannot prevent an *a priori* skeptical challenge from being raised simply by insisting that Descartes' and Wittgenstein's skeptical hypotheses are impossible. The belief that these hypotheses are impossible is an *a priori* belief—i.e., *the very sort of belief that a priori skeptical hypotheses seek to call*

into question. Just as one cannot amass inductive evidence to keep the problem of induction from arising or amass testimonial evidence to prevent the problem of other minds from arising, one cannot appeal to one's *a priori* beliefs about the impossibility of *a priori* skepticism to prevent *a priori* skeptical challenges from being raised.[12]

V.

Call the requirement that it must be possible for skeptical hypotheses to be true and for targets of skeptical attack to be false the 'possibility requirement.' I have argued that the possibility requirement does not demand logical or metaphysical possibility. After briefly considering (in the present section) whether the requisite sort of possibility is epistemic, I argue (in the following section) that subjective indistinguishability—understood as an explicitly modal notion—is all that is required.

Let us begin by noting that we often speak intelligibly about the possible falsehood of certain necessary propositions. Consider, for example, an assertion of the following sentence made by Saul Kripke in 1970:

(11.1) Fermat's Last Theorem might not be true.[13]

Or an assertion of the following made by someone in ancient Rome:

(11.2) Hesperus might not be Phosphorus.

Each of these assertions seems correct, even though we now know that Fermat's Last Theorem is true and that Hesperus is identical to Phosphorus. Our present knowledge, however, does not prevent other speakers who lack this information from saying something true by assertively uttering these sentences. In fact, even if Andrew Wiles's proof of Fermat's Last Theorem is sound, someone who had doubts about certain features of the proof might nonetheless speak truly in uttering (11.1) at the present time. Consider also assertions of the following sentences made by someone today:

(11.3) Goldbach's conjecture might be true.

(11.4) Goldbach's conjecture might be false.

Since our present knowledge is currently unable to decide whether Goldbach's conjecture is true, an assertion of either (11.3) or (11.4) would be correct—even though if the conjecture is true, it is necessarily true, and if false, it is necessarily false. Note that most of the foregoing assertions would be false if they were taken to express logical or metaphysical possibilities. Because these assertions seem to be correct and because their apparent correctness has something to do with what speakers know, they are commonly taken to express epistemic possibilities.[14]

Consider the fact that skeptics try to get non-skeptics to concede the truth of only the first members of the following pairs of propositions[15]:

(12.1) I might be a BIV.

(12.1') I might have been a BIV.

(12.2) *Modus ponens* might not be correct.

(12.2') *Modus ponens* might not have been correct.

(12.3) '2 + 3 = 5' might not be true.

(12.3') '2 + 3 = 5' might not have been true.

In order to know whether any of the second members of these pairs of propositions is correct, we would need to know what kind of world we inhabit. If, for example, I inhabit a "normal" world, (12.1') will be false. But if I inhabit a BIV world, it will be true. Skeptics and non-skeptics, however, can both agree that the first members of each pair of propositions are correct without either party begging the question against the other about our knowledge of the world we inhabit. This suggests that when skeptics assert that we might be brains in vats, deceived by an evil demon, etc., what they might have in mind is the epistemic possibility that we are. Perhaps, then, the following constraints on skeptical hypotheses are correct:

(SH8) In order for a skeptical hypothesis, *SK*, to raise a significant skeptical challenge to *S*'s putative knowledge that *O*, it must be epistemically possible for *O* to be false.[16]

(SH9) In order for a skeptical hypothesis, *SK*, to raise a significant skeptical challenge to *S*'s putative knowledge that *O*, it must be epistemically possible for *SK* to be true.

Unfortunately, however, epistemic possibility is commonly understood in ways that cannot help to explain the possibility requirement.

Consider what is perhaps the most common definition of epistemic possibility:

(13.1) *p* is an epistemic possibility for *S* iff *p* is compatible with what *S* knows.[17]

The central problem with (13.1) for our purposes is that if I in fact know that I have hands, it will not be epistemically possible that I am being deceived into falsely believing that I have hands. Thus, if (SH9) were true, no skeptical challenge could ever be raised against a belief that counts as knowledge. Yet this is clearly implausible. Even if I know that I have hands, epistemological challenges can still be raised against this belief. My knowledge of this fact may aid me in responding to such challenges, but it cannot prevent those challenges from ever being raised. Furthermore, whenever we acknowledge that a skeptical argument has raised a significant challenge and we admit that it is in some sense an open question whether we know the things we think we know, we are not necessarily admitting that being a handless brain in a vat really is compatible with what we know about the world. Non-skeptics can take their favored responses to skepticism to show that a skeptical hypotheses can be known to be false and yet can admit that it raises a significant skeptical challenge. The sense in which it is an open question whether are brains in vats, deceived by an evil demon, etc., then, cannot be cashed out in the terms provided by (13.1).

Similar difficulties beset other definitions of epistemic possibility, such as the following[18]:

(13.2) *p* is an epistemic possibility for *S* iff *S* does not know that *p* is false.[19]

(13.3) *p* is epistemically possible (for the relevant community) iff no one in the relevant community knows that *p* is false and there is no practicable investigation by means of which members of the relevant community could establish that *p* is false.[20]

(13.4) *p* is epistemically possible (for the relevant community) iff *p* is not known to be false by any member of community *C*, nor is there any member of *C* such that if that person were to know all the propositions known to members of *C*, she could on the basis of this knowledge come to know that *p* is false.[21]

(13.5) *p* is epistemically possible for *S* iff nothing that *S* knows entails, in a manner obvious to *S*, not-*p*.[22]

(13.6) *p* is epistemically possible for *S* iff not-*p* is neither taken by *S* as known nor can be recognized to be metaphysically impossible *a priori* (i.e., regardless of the particular state of information *S* is in).[23]

(13.7) *p* is epistemically possible for *S* iff *p* is true or *S* does not have justification for not-*p* adequate for dismissing *p* or *S*'s justification for not-*p* is not Gettier-proof.[24]

(13.8) A use of 'It is possible that *p*' is true in a context of assessment (i.e., a setting in which such a use is being assessed for truth or falsity) iff the proposition expressed by the use is not ruled out by what the subject(s) making the assessment know(s).[25]

(13.9) *p* is epistemically possible iff *p* cannot be ruled out by *a priori* reasoning.[26]

The combination of any of these views—with the exception of (13.9)—with (SH8) and (SH9) would imply that no skeptical challenge could be raised against any belief that counts as knowledge. And although most skeptical hypotheses that concern our knowledge of the external world would satisfy (13.9), this definition of epistemic possibility cannot help us understand the sense in which *a priori* skeptical hypotheses are possible. Thus, however useful these notions of possibility may be for other theoretical purposes, they cannot help us understand the possibility requirement. We must, then, look elsewhere for an account of the sense in which skeptical hypotheses must be possibly true and the targets of skeptical attack possibly false.[27]

VI.

I propose that the correct way to understand the possibility requirement is simply in terms of subjective indistinguishability. As a first pass, consider David Lewis's (1996, 552-553) definition of possibilities left uneliminated by a subject's evidence (where evidence is not taken to be equivalent to what a subject knows):

(14.1) *W* is a possibility for *S* iff in *W S*'s perceptual experiences and memories match *S*'s perceptual experiences and memories in actuality.

Lewis (1996, 553) explains that the relevant possibilities are "those in which the subject's entire perceptual experience and memory are just as they actually are." If there is such a thing as narrow mental content, it may be that *W* is possible for *S* just when *S*'s narrow contents in *W* match *S*'s narrow contents in actuality.[28] The possibilities in question are not simply possibilities as to how the entire world is. Rather, they are possibilities *de se et nunc*—i.e., possibilities centered on particular subjects. Centered worlds can be thought of as pairs of worlds and designated inhabitants thereof (Lewis 1979, 532), pairs of worlds and viewpoints of particular subjects (Chalmers 1996, 60), or triples of worlds, times and agents (MacFarlane forthcoming). Distinct subjects in the same possible world thus inhabit different centered worlds.

Because Lewis (1996, 553) allows that other forms of basic evidence could be included in (14.1) and because we want to allow for the possibility of *a priori* skeptical hypotheses, we need to consider the following generalization of (14.1), which ranges over all of *S*'s experiences (perceptual, intuitive or otherwise):

(14.2) *W* is a possibility for *S* iff in *W S*'s experiences and memories match *S*'s experiences and memories in actuality.[29]

Call the sense of possibility determined by (14.2) 'experiential possibility.' Experiential possibility is basically subjective indistinguishability. Lewis no doubt intended to limit the possibilities in (14.1) to logical and metaphysical possibilities. However, since I want to use the notion of experiential possibility to

explicate the possibility requirement and because I have argued that skeptical hypotheses can be logically and metaphysically impossible, the domain of (14.2) should be understood to include logical and metaphysical impossibilities.

I contend that the following constraints jointly constitute the possibility requirement on skeptical hypotheses:

(SH10) In order for a skeptical hypothesis, *SK*, to raise a significant skeptical challenge to *S*'s putative knowledge that *O*, it must be experientially possible for *O* to be false.

(SH11) In order for a skeptical hypothesis, *SK*, to raise a significant skeptical challenge to *S*'s putative knowledge that *O*, it must be experientially possible for *SK* to be true.

Thus, the central feature of those possibilities in which we are brains in vats or deceived by an evil demon is that they are experientially possible. The inductive skeptic's worlds in which the future fails to resemble the past may be nomologically or even epistemically impossible, but they are nonetheless experientially possible. And although the falsity of certain classes of necessary truths is neither logically, metaphysically nor (in many cases) epistemically possible, the experiential possibility of their falsity is enough to give the *a priori* skeptic an epistemological foothold from which to lodge a significant skeptical challenge.

The motivation for using possibilities that are centered on particular subjects arises naturally from the way that skeptical challenges—particularly those concerning the external world—have traditionally been raised. The external world for me includes, among other things, the Empire State Building, Mt. Everest and your mind. But the external world for you (presuming you exist) includes my mind but not yours. It is commonly assumed that a successful reply to skepticism should appeal only to resources available from reflection on one's own thoughts without supposing that anything besides one's own consciousness exists.[30] The egocentric nature of the traditional skeptical challenge is thus reflected in the subjective form of possibility to which (SH10) and (SH11) appeal.

Let '*U*' denote the set of possible and impossible (uncentered) worlds in which I exist, and let '*V*' denote the set of centered worlds that result from taking each member of *U* and centering it on me. Let '*W*' denote the set of worlds within *V* in which I fail to have any knowledge of a particular kind, and let '*X*' denote the set of worlds within *V* that are subjectively indistinguishable from the actual world. *X*, of course, can be subdivided into the set of *X*-worlds that are also *W*-worlds and the set of *X*-worlds that are not *W*-worlds. Call the former set '*Y*' and the latter '*Z*.' The (centered) actual world will be a member of either *Y* or *Z* but not both. The heart of the skeptical challenge, then, is this: the skeptic alleges is that my evidence is insufficient to tell me whether I inhabit a *Y*-world or a *Z*-world. Because skeptical challenges that are directed to you will involve different sets of centered worlds, skeptical challenges can be seen to be personalized in a certain sense.

A brief word about impossible worlds may be in order before drawing things to a close. The notion of experiential possibility brings with it a commitment to metaphysically impossible worlds.[31] How costly is that theoretical commitment? I suggest it is not costly at all. Impossible worlds are already needed for making sense of counterpossibles (Mares 1997; Nolan 1997; Vander Laan 2004), the propositional content of necessarily false beliefs (Restall 1997), paraconsistent logic, relevant logics, and indeed any alternative logic that may be incorrect. Furthermore, on the vast majority of conceptions of possible worlds, a commitment to impossible worlds is entirely innocent. Extreme modal realists (e.g., Lewis 1986) will obviously have difficulty accommodating impossible worlds.[32] But perhaps this is simply one more reason to reject their conception of worlds. Ersatzists (e.g., Carnap 1947; Jeffrey 1965; Adams 1974; Plantinga 1974; 1976; 1987; Lycan 2002) who take possible worlds to consist in sets of sentences or propositions can simply take impossible worlds to be inconsistent sets of sentences or propositions.[33] Modal fictionalists (e.g., Rosen 1990; Nolt 1986) who deny that possible worlds talk brings with it any ontological commitment to possible worlds should equally have no difficulty adopting impossible worlds. If possible worlds are merely the products of fiction, it seems we could simply tell our fictional story about

worlds in a way that includes impossible worlds. Combinatorial theorists (e.g., Skyrms 1981; Armstrong 1986; 1989) who take possible worlds to be recombinations of the fundamental elements of the actual world can simply take impossible worlds to be impossible recombinations. Daniel Nolan (2002, §1.5) even suggests that non-fictionalists about possible worlds could be fictionalists about impossible worlds. Consequently, the fact that my account of the possibility requirement brings with it a commitment to impossible worlds (or at least to impossible worlds talk) should not be taken to be a theoretical liability for the view.

VII.

I have argued that in order for skeptical hypotheses to be used to raise significant skeptical challenges, they do not need to be: (i) incompatible with targets of skeptical attack, (ii) logically or metaphysically possible, (iii) epistemically possible or (iv) plausible. I have also argued that in order for our ordinary beliefs to serve as the targets of significant skeptical attacks, (i) these beliefs do not need to be *a posteriori* and (ii) that their falsity does not need to be logically, metaphysically or epistemically possible. Instead, I have argued (i) that the truth of skeptical hypotheses and the falsity of targets of skeptical attack must be experientially possible and (ii) that skeptical hypotheses must indicate how subjects can believe what they do on the basis of their evidence while failing to have knowledge. Not only are these conditions necessary for using skeptical hypotheses to lodge effective skeptical challenges, they also seem to be sufficient for doing so. I hope that by reflecting on the proper constraints on skeptical hypotheses I have not only illuminated the nature of skeptical challenges themselves but also broadened our conception of the legitimate kinds of skeptical challenges that can be raised.[34]

References

[Author's Publication]

Adams, Robert. 1974. "Theories of Actuality." *Noûs* 8:211-31.

Armstrong, D. M. 1986. "The Nature of Possibility." *Canadian Journal of Philosophy* 16:575-594.

_____. 1989. *A Combinatorial Theory of Possibility*. Cambridge: Cambridge University Press.

Bealer, George. 2004. "The Origins of Modal Error." *Dialectica* 58:11-42.

Brogaard, Berit & Joe Salerno. 2007. "Why Counterpossibles are Non-Trivial." *The Reasoner* 1:5-6.

Brueckner, Anthony. 1994. "The Structure of the Skeptical Argument." *Philosophy and Phenomenological Research* 54:827-835.

Carnap, Rudolf. 1947. *Meaning and Necessity*. Chicago: University of Chicago Press.

Chalmers, David J. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press.

_____. 2002a. "The Components of Content." In David J. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press, pp. 608-633.

_____. 2002b. "On Sense and Intension." *Philosophical Perspectives* 16:135-182.

_____. (forthcoming). "The Nature of Epistemic Space." Available at: http://consc.net/papers/espace.html.

Cohen, Stewart. 1998. "Two Kinds of Skeptical Argument." *Philosophy and Phenomenological Research* 58:143-159.

Davies, Martin & Lloyd Humberstone. 1980. "Two Notions of Necessity." *Philosophical Studies* 38:1-30.

DeRose, Keith. 1991. "Epistemic Possibilities." *Philosophical Review* 100:581-605.

_____. 1998. "Simple 'Might's, Indicative Possibilities and the Open Future." *Philosophical Quarterly* 48:67-82.

_____. 1999a. "Can It Be That It Would Have Been Even Though It Might Not Have Been?" *Philosophical Perspectives* 13:385-413.

_____. 1999b. "Introduction: Responding to Skepticism." In Keith DeRose & Ted. A. Warfield (eds.), *Skepticism: A Contemporary Reader*. Oxford: Oxford University Press, pp. 1-24.

Descartes, René. 1984. *The Philosophical Writings of Descartes*. Trans. by John Cottingham, Robert Stoothoff and Dugald Murdoch. Vols. 1 & 2. Cambridge: Cambridge University Press.

Edgington, Dorothy. 2004. "Two Kinds of Possibility." *Proceedings of the Aristotelian Society* (Suppl.) 78:1-22.

Egan, Andy, John Hawthorne & Brian Weatherson. 2005. "Epistemic Modals in Context." In Gerhard Preyer & Georg Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. New York: Oxford University Press, pp. 131-169.

Forrest, Peter. 1986. "Ways Worlds Could Be." *Australasian Journal of Philosophy* 64:15-24.

Freud, Sigmund. 1927/1961. *The Future of an Illusion*, trans. James Strachey. New York: W. N. Norton.

Hacking, Ian. 1967. "Possibility." *Philosophical Review* 76:143-168.

Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hetherington, Stephen Cade. 1996. "Gettieristic Scepticism." *Australasian Journal of Philosophy* 74:83-97.

Hintikka, Jaakko. 1962. *Knowledge and Belief*. Ithaca: Cornell University Press.

Huemer, Michael. 2007. "Epistemic Possibility." *Synthese* 156:119-142.

Jeffrey, Richard. 1965. *The Logic of Decision*. New York: McGraw-Hill.

Kripke, Saul. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.

Levin, Michael. 2000. "Demons, Possibility and Evidence." *Noûs* 34:422-440.

Lewis, David. 1979. "Attitudes *De Dicto* and *De Se*." *Philosophical Review* 88:513-543.

_____. 1986. *On the Plurality of Worlds*. Oxford: Basil Blackwell.

_____. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74:549-67.

Lycan, William G. 2002. "The Metaphysics of Possibilia." In Richard M. Gale (ed.), *The Blackwell Guide to Metaphysics*. Malden, MA: Blackwell, pp. 303-316.

MacFarlane, John. forthcoming. "Epistemic Modals are Assessment-Sensitive."

Mares, Edwin D. 1997. "Who's Afraid of Impossible Worlds?" *Notre Dame Journal of Formal Logic* 38:516-526.

Moore, G. E. 1959. "Certainty." In *Philosophical Papers*. London: George Allen & Unwin, pp. 227-251.

_____. 1962. *Commonplace Book: 1919-1953*, ed. Casimir Lewy. London: George Allen & Unwin.

Nolan, Daniel. 1997. "Impossible Worlds: A Modest Approach." *Notre Dame Journal of Formal Logic* 38:535-572.

_____. 2002. "Modal Fictionalism." In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Available at: http://plato.stanford.edu/entries/fictionalism-modal/index.html.

Nolt, John E. 1986. "What Are Possible Worlds?" *Mind* 95:432-445.

Plantinga, Alvin. 1974. *The Nature of Necessity*. Oxford: Clarendon Press.

_____. 1976. "Actualism and Possible Worlds." *Theoria* 42:139-60.

_____. 1987. "Two Concepts of Modality: Modal Realism and Modal Reductionism." *Philosophical Perspectives* 1:189-231.

_____. 1993. *Warrant and Proper Function*. Oxford: Oxford University Press.

_____. 2000. *Warranted Christian Belief*. Oxford: Oxford University Press.

Pritchard, Duncan. 2005. "The Structure of Sceptical Arguments." *Philosophical Quarterly* 55:37-52.

Rosen, Gideon. 1990. "Modal Fictionalism." *Mind* 99:327-354.

Skyrms, Brian. 1981. "Tractarian Nominalism." *Philosophical Studies* 40:199-206.

Sosa, Ernest. 1999. "Skepticism and the Internal/External Divide." In John Greco & Ernest Sosa (eds.), *The Blackwell Guide to Epistemology*. Malden, MA: Blackwell, pp. 145-157.

Stalnaker, Robert. 1984. *Inquiry*. Cambridge, MA: MIT Press.

Stanley, Jason. 2005. "Fallibilism and Concessive Knowledge Attributions." *Analysis* 65:126-131.

Steup, Matthias. 2005. "Knowledge and Skepticism: Introduction." In Matthias Steup & Ernest Sosa (eds.), *Contemporary Debates in Epistemology*. Malden, MA: Blackwell, pp. 1-12.

Stroud, Barry. 1965. "Wittgenstein and Logical Necessity." *The Philosophical Review* 74:504-518.

_____. 1984. *The Significance of Philosophical Skepticism*. Oxford: Clarendon Press.

Teller, Paul. 1972. "Epistemic Possibility." *Philosophia* 2:303-320.

Van Cleve, James. 1979. "Foundationalism, Epistemic Principles, and the Cartesian Circle." *Philosophical Review* 88:55-91.

Vander Laan, David. 2004. "Counterpossibles and Similarity." In Frank Jackson & Graham Priest (eds.), *Lewisian Themes: The Philosophy of David K. Lewis*. Oxford: Oxford University Press, pp. 258-275.

Vogel, Jonathan. 2004. "Skeptical Arguments." *Philosophical Issues* 14:426-455.

Wittgenstein, Ludwig. 1953. *Philosophical Investigations*, 3rd edn. Trans. by G. E. M. Anscombe. New York: Macmillan.

_____. 1956. *Remarks on the Foundations of Mathematics*. Edited by G. H. von Wright, R. Rhees & G. E. M. Anscombe. Trans. by G. E. M. Anscombe. Oxford: Oxford University Press.

Yagisawa, Takashi. 1988. "Beyond Possible Worlds." *Philosophical Studies* 53:175-204.

[1] Adapted from Pritchard (2005, 39)

[2] Formulating the precise epistemic principle(s) upon which (3.2) and (5.1) are based can be a tricky affair. Some philosophers (e.g., Sosa 1999, 145; Stroud 1984) have suggested the following:

(PE) If *S* knows that *p* and *S* knows that *q* is incompatible with *S*'s knowing that *p*, then *S* must be able to rule out *q*.

An apparently serious problem with (PE), however, is that it entails the KK-principle. Since my failing to know *p* is incompatible with my knowing that *p*, (PE) implies that I must know that I do not fail to know *p* in order to know that *p*. In other words, in order to know, I must know that I know. Cf. Vogel (2004) for further discussion of this issue.

[3] Note that, given closure, if Fred is to know that the animal in the pen is a zebra and Fred knows that the animal in the pen's being a zebra is incompatible with the animal in the pen's being a lion, then Fred needs to be in a position to know that the animal in the pen is not a lion. Veridical substitutability into the closure principle schema, then, is not sufficient for giving a proposition any skeptical force.

[4] I am indebted to [AC] and [KE] for helpful suggestions on how to formulate this principle.

[5] The following potential problem for (SH3) was raised by [KE]: Suppose *S* believes the animal in the pen is a zebra but is unable to tell zebras from lions. This supposition seems to satisfy (SH3), yet it does not seem like a skeptical hypothesis. I think the solution to the difficulty is to note that the *O*'s we are dealing with are restricted to propositions we are already ordinarily inclined to take subjects to know. Since we would not be inclined in the described case to think the subject knows, it wouldn't count as a relevant *O*.

[6] Thanks to [AC] and [KE] for helpful comments on the proper reformulation this principle.

[7] Thanks to [AT] for helping me avoid an important error concerning the scope of this *a priori* skeptical challenge.

[8] Cf. Hetherington (1996) for a defense of the view that Gettier cases can be used as skeptical hypotheses.

[9] Cf. section VI below for the requisite explication.

[10] For various reasons, ordinary closure and underdetermination principles cannot be used to underwrite this premise. I hope to address these difficulties in subsequent work.

[11] Like an increasing number of contemporary philosophers (e.g., Mares 1997; Nolan 1997; Vander Laan 2004; Brogaard & Salerno 2007), I do not think that subjunctive conditionals with impossible antecedents (i.e., counterpossibles) are vacuously true.

[12] Some philosophers contend that skeptical arguments based upon necessarily false skeptical hypotheses will always be epistemically self-defeating. Cf. [author's publication] for further discussion of this charge.

[13] Cf. Kripke (1980, 141). Fermat's Last Theorem was not proven until 1995.

[14] Cf. DeRose (1991; 1998; 1999a) for a detailed defense of this view.

[15] The first members of each pair roughly correspond to what two-dimensional semanticists (e.g., Davies & Humberstone 1980; Chalmers 2002a; 2002b; forthcoming) have called 'possibilities considered as actual,' whereas the second members correspond to 'possibilities considered as counterfactual.'

[16] A further difficulty with (SH8) is that an increasing number of philosophers are coming to the defense of the following epistemic principle:

(KE) If not-*p* is epistemically possible for *S*, then *S* doesn't know that *p*.

DeRose (1991, 593-594), for example, argues that statements of the form 'I know that *p*, but *p* might be false' are not merely false but contradictory. (SH8) and (KE) jointly imply that skeptical challenges can only be raised against propositions we fail to know. Again, that is implausible. Although (KE) is controversial, I do not want my account of the proper constraints on skeptical hypotheses to prejudge the issue here. Cf. Hintikka (1962), Hawthorne (2004), and Huemer (2007) for other defenses of the view.

[17] Cf. Hintikka (1962), Stalnaker (1984, 143) and Hawthorne (2004, 26).

[18] These difficulties also afflict accounts that define epistemic possibility in terms of certainty and take certainty to be necessary for knowledge. Cf. Van Cleve (1979, 63) and Moore (1962, 184). It should be remembered that Moore's *Commonplace Book* is a compilation of notebooks that were never properly prepared for publication.

[19] Hacking (1967) and others argue that (13.2) is also problematic because it counts as epistemically possible propositions that *S* has conclusive evidence against in cases where *S* does not realize *S* possesses this evidence.

[20] Cf. Hacking (1967). DeRose (1991, 593-594) offers a version of this view that allows for wide contextual variability both in what the relevant community is and in what counts as a relevant way of coming to know. This means that the content and truth values of statements of epistemic possibility will vary across conversational contexts. In contexts where low-standards for knowledge are in place, it will be true to assert both 'I know I have hands' and 'It is not epistemically possible that I am a brain in a vat.' In other contexts, however, where my knowledge of the external world is in question, 'It is epistemically possible that I am a brain in a vat' will be true.

[21] Cf. Teller (1972, 310-311). Teller suggests a contextualist treatment of the boundaries of the relevant community. His definition also embeds a counterfactual, which must be cashed out at least partly in terms of metaphysical possibility.

[22] Cf. Stanley (2005, 128).

[23] Cf. Edgington (2004, 6).

[24] Cf. Huemer (2007, 129).

[25] Cf. MacFarlane (forthcoming). Cf. Egan, Hawthorne & Weatherson (2005) for further discussion of this kind of relativism about epistemic modals.

[26] Cf. Chalmers (2002a, 609-610).

[27] The fact that (13.1) and (13.5) make use of logical notions such as entailment or consistency may pose a problem for understanding the sense in which logically impossible skeptical hypotheses represent genuine possibilities for us.

[28] Cf. Levin (2000) for an account of skeptical challenges along these lines.

[29] The centered possibilities represented here may be analogous to Kripke's (1980, 103ff) qualitatively identical epistemic situations.

[30] Cf. Levin (2000, 424) for further discussion of this "old-fashioned constraint" on replies to skepticism.

[31] In an effort to avoid appealing to impossible worlds while allowing certain metaphysical impossibilities to be epistemically possible, Chalmers (2002a; 2002b; forthcoming) avoids representing or associating every maximal epistemic possibility with a world. Instead, he takes a 'scenario' to be a maximally specific way the world might (epistemically) be and associates epistemic possibilities with scenarios. Chalmers (forthcoming, §4) then introduces a technical distinction between 'satisfaction' and 'verification' via the following examples. Our world satisfies the thought expressed by 'water is H_2O.' A Twin Earth world that is superficially like our world but in which the oceans and lakes are filled with clear, drinkable XYZ verifies the thought expressed by 'water is H_2O' but does not satisfy that thought. Chalmers (forthcoming, §4) offers the following test to determine whether a world verifies or satisfies a thought:

> To determine whether the Twin Earth world satisfies "water is XYZ" we can ask: if the liquid in the oceans and lakes *had been* XYZ would water *have been* XYZ? If Kripke and Putnam are correct, the answer is no. Verification, on the other hand is a broadly epistemic notion, concerning ways the world might be. To determine whether the Twin Earth world verifies "water is XYZ," we can ask: if the liquid in the oceans and lakes *is* XYZ, *is* water XYZ?... And in this case,… the answer is yes.

Asking the question about XYZ in the subjunctive corresponds to considering the Twin Earth world as counterfactual, whereas asking the indicative question corresponds to considering the world as actual. Considering a world as counterfactual involves taking the nature of the actual world to be fixed and considering the counterfactual world in relation to these facts.

Chalmers contends that by employing the notions of satisfaction and verification he can define both epistemic and metaphysical possibilities over the same set of metaphysically possible worlds and can avoid defining truth at a scenario. Instead of claiming that a scenario *S* verifies a thought *T* iff *T* is true at *S*, Chalmers employs the ordinary notion of truth at a world and claims that a scenario *S* verifies a thought *T* iff the value of the function verifies(*S*, *T*) is true. This enables Chalmers to claim that the thought expressed by 'water is not H_2O' is epistemically possible for certain subjects, even though there is no world in which water is not H_2O.

[32] Cf. Yagisawa (1988) for an attempt to accommodate impossible worlds within the Lewisian framework.

[33] Because ersatz possible worlds are abstract rather than concrete, there will be no further difficulty in supposing that impossible worlds exist than there is in supposing that possible worlds exist. According to Adams's (1974, 225) actualist account of possible worlds, a possible world or "world story" is a maximally possible set of propositions—a set of propositions such that (i) for any proposition *p*, the set contains either *p* or its negation, not-*p*, and (ii) it is metaphysically possible for all of the members of the set to be true together. Ersatz impossible worlds can be defined in an analogous fashion: an impossible world is a maximal set of propositions—i.e., a set of propositions such that (i) for any proposition *p*, the set contains either *p* or its negation, not-*p*, and (ii) it is not metaphysically possible that all the members of the set be true together. A proposition will be true in a world story just when it is a member of that story. According to Plantinga's (1974; 1976; 1987) actualist account of possible worlds, non-actual possible worlds are maximal states of affairs that do not obtain. Since Plantinga (1976, II.1) believes that impossible states of affairs such as the number 9's being prime do not cease to be genuine states of affairs simply because they could not have obtained, it is but a short step to admit that there are maximal states of affairs (i.e., worlds) that could not have obtained.

Plantinga's (1987, 192) definition of a 'maximal state of affairs' applies only to possible states of affairs: "a possible world is a maximal possible state of affairs, where a state of affairs *S* is maximal if and only if for every state of affairs *S**,either *S* includes *S**or *S* includes the complement ~*S** of *S**." However, Plantinga's (1974, 46) definition of 'the book on *W*,' for each world *W*, is amenable to our current purposes. According to Plantinga, for every proposition *p*, either *p* is a member of the book on *W* or not-*p* is. We can then say that a world is maximal iff for every proposition and every state of affairs, either the state of affairs corresponding to *p* or the state of affairs corresponding to not-*p* is included in *W*.

Linguistic ersatzists (e.g., Jeffrey 1965, 196-197; Carnap 1947) take possible worlds to be maximal sets of sentences, where a set is maximal iff for every sentence, either it or its negation is a member of the set. On this view, an impossible world would be a maximal set of sentences that is not compossible. Forrest (1986, 15) offers an ersatzist account of non-actual worlds

in terms of uninstantiated "world-natures." A world-nature is a complex property which is the conjunction of all of the non-relational properties of what we ordinarily think of as a possible world. Forrest (1986, 21) suggests: "on my theory, you can have impossible 'world-natures' *if you want them.* For example, you could insist that there is a conjunction of *being round, being square, being uniformly vermillion,* and *being uniformly turquoise*."

> Some ersatzists argue strongly in favor of adopting impossible worlds. Lycan (2002, 312), for example, writes:
>
> [The] fact that Lewis rules out impossible worlds… is a serious liability. For linguistic semantics needs impossible worlds. Conditional sentences can have impossible antecedents, as in "If there were round squares,…," and people can often be described as having contradictory beliefs. (Moreover, I can think of no direct argument for "nonexistents" that does not support impossibilia by parity of reasoning; I would not expect anyone to find a reason, independent of Concretism, for countenancing non-actual possible worlds but refusing to acknowledge impossible ones.)

[34] I would like to thank audience members at the 2008 Meeting of […] and participants in my Spring 2008 graduate epistemology seminar for helpful comments and criticisms on earlier versions of this paper.